

1 Introduction

Begin with our title, “Dynamics of Public Opinion.” Why opinion dynamics? Why do we care how public sentiment moves over time? We could care about opinions for the same sort of reasons that collectors care about coins or stamps, just for the joy of collecting. But we do not collect just because we like collecting. (And, in fact, collecting data from thousands of surveys as we have done here does not necessarily spark joy; it is not satisfying on its own.) We care about opinion dynamics because theory – democratic theory – requires that we do. Democratic theory mandates a system in which policy-making responds to public demands. And we cannot understand the black box that is politics unless we understand public demands.

In this Element we aim to get inside that box, called representation, to gain purchase on American politics by coming to understand how opinion moves. That is our goal. We are ambitious. We wish to understand *all* of what Americans want from government, spanning multiple dimensions. We will propose that three models – only three – are needed to understand virtually all opinion movement in American politics. Jointly the three models form a theory of opinion movement. We will observe that public opinion sometimes cycles over time, sometimes trends in one direction, and sometimes (rarely) does not move at all. Our definition of success in this endeavor will require us to take everything we know about politics, and how political parties compete over policy-making and elections, to explain the variations in opinion movement. To the extent that we succeed in these pages, the reader will come to think that the whole system we propose could not be otherwise. But of course success in scholarship is halting and partial. Readers perhaps will be unconvinced.

The study of representation hinges on accurately measuring public preferences for government services, and assessing the relationship between those demands and government outputs. This Element makes advances in both areas. First, we develop a theory that outlines the important role political parties play in facilitating public responsiveness to government action. On issues where the parties regularly take opposing positions (most welfare state policies for example), the public can infer the direction of policy change in Washington simply by knowing which party is in control of the White House. Opinion, therefore, cycles over time in response to changes in party control. In fact, this is the single most prominent mechanism of opinion change, explaining the majority of the cases we will analyze in the pages to come. But, as good as it is, the party-based model is an incomplete one; there are two others. Where party cues are absent, or where cleavages other than partisanship are dominant, different patterns of public response emerge.

One understudied and generally poorly understood area of opinion change relates to those areas where the political parties do not take opposing positions. On some issues, nearly all Americans agree, or their disagreements relate to things other than their party preferences, ideological positions, or other factors connected to elections and voting. Political scientists (and public opinion polling firms) typically pay little attention to these areas of public policy, since by definition they are aside from the stuff of politics; they will not affect any election outcomes. If these issues don't regularly divide the major parties and don't align with the ideological divides that separate us, these issues cannot logically affect electoral outcomes. But they still exist; people still have those attitudes, and we should understand them. We make some initial steps toward doing so here. They may cycle, or remain relatively stable over time, but they cannot cycle in the same way that partisan issues cycle. This is because the main driver of change for partisan issues – citizen response to the inferred direction of policy change based on control of the White House – is impossible. If the parties don't differ in their positions on these issues, then the public cannot infer anything related to them on the basis of which party controls the government. So we will develop one model for issues that regularly divide the parties, one for those that don't, and a third, which follows.

Our third and last model of attitude change over time relates to issues that concern long-term cultural shifts. These are social transformations affecting society in powerful ways, literally shifting the norms of cultural acceptability of a given issue position. These can be so powerful that they overwhelm the influence of any short-term partisan differences, driving substantial shifts in public opinion over time, all in the same direction. We use the examples of racial attitudes, attitudes toward women's equality, and attitudes toward the rights of gay Americans to illustrate this third pattern of change. Change in these "cultural shift" policy domains is steady over the long term. The influence of short term partisan differences on these issues is overwhelmed by the longer-term trends associated with two factors: (1) large swaths of the American public progressively adopting new, pro-equality positions on the issue, and (2) the generational replacement of individuals with once-widespread but no-longer-majority anti-equality opinions with younger individuals coming of age during a different time, and reflecting more progressive positions on these cultural shift issues.

The first two models of opinion dynamics we identify imply no long-term change in opinion at all. In these cases, aggregate opinion moves up and down (or, left and right) but fifty years later remains roughly where it started. Cycles and movements occur, but these are centered around some mean opinion, and this mean opinion shows relative stability when considered over a long period

of time. The third model is not like that at all. Absolute change associated with issues of cultural shift reflects shifting societal values and results in long-term opinion trends.

Our simple proposal is that there are three, and only three, models needed to understand opinion dynamics. These relate to three types of issues: partisan, nonpartisan, and cultural shift. Partisan issues cover the bulk of public opinion questions of interest to most political scientists, and most often posed by commercial and academic polling organizations. These correspond to what has been called the “thermostatic” model of opinion change (Soroka and Wlezien 2010; Wlezien 1995). Nonpartisan issues will fail any statistical test for thermostatic response to change in party control of government because the key driver there is missing: there is no difference in expected policy outcomes on these issues based on who is in control. Both parties support the maintenance of our national parks, for example; we cannot infer a position on the space program based on which party is in power. So, if opinion to these policies remains stable, drifts, or cycles, it will do so out of sync with partisan issues. Finally, for cultural shift issues, trends will occur. Absolute opinion change, combined with generational replacement, facilitates long-term absolute shifts in aggregate public opinion. These are so powerful that they overwhelm any short-term partisan differences that may surface. No political party can ignore these trends for long and remain in power – public policies change dramatically over time in response to shifting cultural values, which in turn reinforces the opinion trends.

Our starting point for the theoretical framework we develop is the well-known thermostatic model. We support this model and believe it is simple, accurate, and useful. We build on it by outlining the key role of consistent party cues in shaping citizen inferences. Consistent party cues allow citizens to draw inferences about the direction of public policy simply by knowing which party controls the levers of government. The simplicity of this model is one of its merits, and it also defines the limits of its applicability: to all issues where the parties regularly take opposing sides. That is, logically the model should hold only when such cues are present. Where consistent party cues are not present, the model cannot hold. And, where cultural shifts are so powerful as to overwhelm partisan differences, the model cannot hold. So, we start with the thermostatic model and add two adjustments. The resulting tri-fold model of opinion change incorporates the original model but completes it with two additional models relevant to different types of issues.

None of what we write here should diminish the power of the thermostatic model. In fact, because most policy issues are consistently politicized – and because we show that very low levels of public knowledge about government are needed for the thermostat to work – the model applies to the vast majority of

cases. But there are topics on which the parties do not regularly take opposing views. On issues like funding for NASA, social security and drug treatment programs, the parties typically agree. Policy outputs in these domains remain consistent over time, even when party control in Washington changes. As a result, public opinion on these topics does not respond to partisan control of government.

Still other issues, like civil rights for racial minorities, women and members of the LGBTQ community, activate cleavages that transcend the partisan divide. Here, social movements have challenged societal norms and promoted the expansion of rights for groups that historically have been marginalized. As individuals embrace these new, inclusive values, demand for government policies that facilitate equality increases. This represents absolute rather than relative opinion change. Public opinion on these issues trends over time rather than cycling in response to changes in party control. These trends are due, in large part, to generational replacement. Individuals from younger cohorts are socialized in a more inclusive environment, leading them to more fully embrace pro-rights values than did members of older cohorts. As public opinion moves in a liberal, pro-equality direction over time, so too, does public policy. And, as we have mentioned, some cultural shifts are so great that, over two decades or more, they bring us to places not previously imagined or taken seriously, such as the legalization of gay marriage.

We begin with the reigning model of the process, the “thermostatic” model of Wlezien and Soroka (Soroka and Wlezien 2010; Wlezien 1995). The thermostatic model uses the metaphor of the home thermostat to explain how citizens regulate policy-making. Its essence is that citizens have a preference for comfortable temperatures. And so when the house strays in one direction or the other from that ideal they can regulate the system simply by turning the thermostat up or down. Note how undemanding the thermostat is. People need not understand the physics of home heating and need not even know the current inside temperature. All they need to know is “too hot” or “too cold” and the corrective action is obvious. Similarly when policy-making is too liberal or too conservative (or perceived to be) they can call for it to reverse to get back to the citizen’s ideal. The general prediction that emerges from this process is that the direction of opinion change is opposite to the dominant direction of change in policy. Liberal policy produces a conservative reaction and conservative policy produces a liberal reaction.

Although we will offer small modifications to the model, we regard it as spectacularly successful. So our critiques to come should be seen as small modifications to a theory we admire. It could not be as successful as it is if it were far wrong.

1.1 The Thermostat and the Implied Thermostat

The thermostatic model of Wlezien (1995) assumes that citizens react to actual policy making. Citizens compare their personal preferences for the level of policy to what government is doing currently and react in the opposite direction if policy is tending too far left or right of their personal ideal point. This model has one very big virtue: it works. Knowing the direction that policy is tending, one can predict that opinion will tend in the opposite direction and be right a good proportion of the time.

But it is a bit strange that the thermostatic model works – and, indeed, works well. What is strange is that the assumption that citizens are knowledgeable of current policy-making is strikingly contrary to most public opinion scholarship. The dominant view of political science is that citizens are normally inattentive, with public affairs well down the list of things they care about, and well down the list of things to which they assign their precious time. Since knowing the current direction of policy is intellectually demanding, it seems unlikely that it could be true that citizens track policy direction for more than the most publicized issues of the moment. When the George W. Bush Administration slashed taxes in 2001 and again in 2003, the public was probably aware that tax policy was moving in the conservative direction. But policy is moving in dozens of domains (many of which we will quantify) and most of these, most of the time, are invisible to all but the hyper attentive.

So how can it be that a theory successfully predicts opinion movement when its central underlying assumption is dubious? Posing that question leads us to postulate a supplemental, and simplifying, mechanism that facilitates knowledge of policy change. We begin an answer by assuming that the mass public in the aggregate knows two facts about American politics. The first is that the Democrats are the party of expanding government liberalism, the party that does more, spends more, and taxes more, while Republicans are and do the opposite.¹ Not everybody knows this, as we have known at least since Campbell (1960), but it is known in the aggregate. Secondly, almost everybody knows who the current president is and which party he or she represents. Knowing those two facts permits an *inference* about the current direction of policy change. If a Democrat occupies the White House, policy is probably changing (in all domains) in a liberal direction. If a Republican is in office, policy change will be toward conservatism. This inference will usually be correct.

¹ Except in the instance of national defense, where Republicans do more and spend more vis-à-vis Democrats.

1.1.1 Party Control and Public Inference

Although short on political knowledge of specific kinds, the electorate does know which party controls the White House at any given time. We postulate that it uses this knowledge to make inferences about public policy. If, for example, Democrats are in power, regardless of the specifics of legislation proposed or acted upon, the public infers that policy is changing in a liberal, government expanding, direction. This is a critical shortcut that reaches essentially similar conclusions about what government is doing (and spending) to what would emerge from detailed and careful knowledge, but obviates the need for attention to details. How would we explain, for example, that Barack Obama's Affordable Care Act was perceived as a decisively liberal approach to health care, when its central feature, the individual mandate, was a product of the conservative Heritage Foundation and the whole plan, based on private sector providers and insurers, emulated the program put forward by Mitt Romney, a Republican, in Massachusetts? The answer is inference. The plan was crafted by President Obama's congressional allies, championed by him, and passed almost entirely by Democratic votes. Therefore, it was liberal by inference.²

In contrast to Soroka and Wlezien (2010), we do not posit that the electorate knows – or needs to know – things like the general level of appropriations. That is a lot of detail and complexity. One can just infer that spending is probably up if Democrats are in the White House or down if Republicans are in control.³ Crucially that means that all that needs to be known with any certainty is which party is in power. The rest can be inferred, with a decent, if imperfect, level of accuracy. Thus we propose opinion that behaves thermostatically, becoming more conservative when liberals are in power and more liberal when conservatives are in power. Therefore our modification of the Wlezien (1995) and Soroka and Wlezien (2010) theory is this: we posit that citizens either know the direction of policy change (probably not very often) or infer it from knowing who occupies the White House and the policy brand of that party (much more often). Since the inference is usually correct, the connection between known or inferred policy change and opinion outcome is robust. This we call the *implied*

² This inference was made stronger still by the policy's popular moniker, "Obamacare." Congressional Republicans, who popularized the label, understood very clearly that linking the policy to the Democratic president would increase public perception that the plan was a liberal one.

³ Is it safe to assume that citizens know which party advocates more spending? It is. For this point data are available. The question has been posed pretty directly in American National Election Studies (ANES) surveys since the 1970s. To cite one example, in the most recent study the modal response on the ANES spending and services tradeoff scale is the most extreme spending category (1) for the Democratic candidate (Biden) and the most extreme "fewer services" category (7) for the Republican candidate (Trump). The scale means are 5.32 (Democratic) and 2.84 (Republican), where 4 is neutral.

thermostatic model. With this, we hope to distinguish our revised model from the original. It amounts to the idea that citizens respond to changes in partisan control of the White House. It is clear that we are pushing the limits of simplicity here to their maximum.

Our preference for simplicity and theoretical clarity is so strong that we also do not incorporate divided versus unified government into our theory, though doing so could be a useful exercise in the future. Clearly, our theory applies most strongly during periods of unified government. In this situation, the public can most confidently infer that one party is in control. And of course, to the extent that congressional majorities are greater, this would be even more strongly the case. In periods of divided control, members of the public may be subjected to mixed messages about who really is “in control” of the government. Still, we expect members of the public to respond to White House control, not control of the House or the Senate. President Trump faced a Democratic House of Representatives during his last two years in office, but the public responded to Trump, not House Speaker Nancy Pelosi in assessing the direction in which the government in Washington was headed. The signal under unified government should be stronger than that under divided government, but even under divided government we expect the public to respond to White House control. Under divided government, it is clear that signals would not be as strong and therefore the thermostatic response might be expected to be weaker. If we were to limit the tests that we conduct in the pages below only to periods of unified government, we would expect them to be even stronger than we show. If we find support in the pages to come for our simple theory, as we do, then a slightly more complicated theory incorporating divided versus unified government, might find even stronger support. But our theoretical interest is not in divided government, so we do not focus on this question. Rather, we want to distinguish among three types of issues, only one of which is subjected to the thermostatic response.

1.1.2 *Giving Credit Where It Is Due*

We did not develop or name the thermostatic model. Wlezien did that. And Soroka and Wlezien refined it. And yet the empirical test we will propose for our version of the model is not the original; indeed we are simplifying in a way that the original developers of the thermostat did not, and we doubt that they made their decisions lightly. To differentiate our simplified model from the original Wlezien “thermostatic model,” we call it the *implied* thermostatic model. Our implied thermostatic model applies widely; in fact it explains the bulk of the opinion series we analyze in the pages to come. But its scope is not universal.

We can now state the limits of the implied thermostat. It works where party cues are regular and powerful, in cases of issue conflict. On issues where the party positions are mixed, confused, or episodic, it does not work. The public cannot infer party positions where the parties do not regularly take opposing positions. There are many issues where the parties take no stand at all. And there are many where the stands require more knowledge of the history of policy disputes than the mass public possesses. Where knowing party positions is easy, the thermostat dominates. Where it is not easy, the thermostat fails.

Our micro theory implies a model of macro-level opinion dynamics. We take that up now.

1.2 Three Kinds of Issues

Begin with partisan issues. These are the main line of party disagreement resulting from The New Deal and Great Society programs, issues about the role of government in the economy, regulation, the social safety net and, in general, the scope of government (how much it should do, how much it should spend, and how much it should tax to pay for that spending). The Democrats regularly take the liberal, government expanding, position. The Republicans regularly take the conservative, government contracting, position. This description fits the general tendency, but the scope of these conflicts has an elastic boundary and changes over time.

Because these sorts of issues are of central importance for party and ideological controversy, they are well and frequently measured in survey research. Opinion surveys tend to measure that which is interesting, and being at the center of party politics makes these issues interesting. Although we have no theory of the potential issue space, we can imagine that it could be vastly larger than these party conflict issues, but with large numbers of *potential* issues lacking the interest that would make them subject to survey questions.

Not every public issue leads to party conflict. Among the policy choices that government must attend to, there are many that may never become party-defining issues. The Federal Government regularly makes decisions about what to do in space travel, how much to fund science, how to regulate public lands, and what rules should govern national parks. In these and many other issue areas the conflicts about what to do have not typically broken along party lines.⁴ As a consequence the distinctive party cues that the thermostatic model requires

⁴ Certainly, there have been brief periods when many of these issues have become politicized at the national level; that is, we might say they have been “episodically” partisan. Recent debates over opening federal land to drilling, for instance, have politicized some aspects of public land management. But in the larger scheme of things – over the span of forty, fifty, or sixty years – policymaking in these areas has not typically been the subject of partisan debate.

do not exist. Democratic and Republican governments make pretty much the same decisions in these areas. Surely there are genuine conflicts in each of these cases. But they are not *party* conflicts. The thermostatic model is not false for these cases. It simply doesn't apply.

Most opinion change in the thermostatic account is relative change. Citizens with relatively fixed views encounter the changing stimulus of the parties alternating in power, and adjust their views accordingly. Liberal party cues make the electorate more conservative and conservative party cues make it more liberal. And this could be true even if no citizen ever changed views in an absolute sense.

The thermostatic model, both in its original formulation and in our slightly altered “implied” version, explains *relative* change in public opinion. Citizens come to support *more* or *less* government *relative* to the level of (perceived) current policy. They could do so with perfectly fixed opinions because the causal force of the model is change in government policy. A citizen who supports level x of policy can rationally call for more or less government policy if the government is currently pursuing more than x or less than x .

But we can imagine what we shall call absolute opinion change as yet a third possibility. In this view, citizens may change their policy preferences relative to their previous preferences, becoming more liberal or conservative on some issue, for example. And the distinctive difference of such absolute opinion change is that it does not depend upon party cues. If something happens that leads people to change their views, some real response to a real stimulus (e.g., war, depression, social movement), then that change will tend to be permanent and it will not be responsive to party cues. And contrary to Soroka and Wlezien's contention, these changes can occur in domains of true public importance. Because we do not posit the party cue as the cause of the change, there is equally no reason to expect a reversal when there is a change of party control.⁵

On this set of potential issues the implied thermostatic model is the wrong model. It posits changing party cues (from changing party control) as the causal force, and for these hypothetical cases (which we will see are real cases) that is simply the wrong causal model. Something else is causing opinion change and it does not cease or alter with a change of party control.

Our theory of the role of party cues thus borrows from the thermostatic account for the important set of issues where the parties regularly offer up

⁵ There is a partially related distinction between absolute and relative in the format of survey questions. Sometimes they ask the absolute in the form “What should government do?” and sometimes they ask the relative, “Should government do more or less than it is currently doing?”

opposite cues. But it also needs to explain opinion dynamics in cases where the party cues do not particularly matter. Thus, party cues play one leading role in our theory of what drives opinion change in the American context.

Of our distinction of issue types, here we deal with the crucially important difference between absolute and relative (i.e., thermostatic) change. We expand this theory of opinion change to include the stimuli of generational changes, our second leading actor. We posit that those stimuli – whether they be the reframing of an issue by a successful social movement, or a set of powerful, connected events – are the force that induces real opinion change among individuals *and* generational change. These stimuli induce absolute opinion change.

1.2.1 Absolute and Thermostatic Opinion Change

Public opinion is in part response to the stimulus of government action. Thus, it may be interpreted as the answer to the question, “What does the public want from government?” Or, since government action changes over time, it may be a relative response to what government has done recently. If policy were as simple as how much to spend, for example, the first type of opinion – which we will call “absolute” – would be the dollar figure for a particular program. The second type – “relative” or thermostatic – would express a desire for more or less spending than the current level.

Relative attitudes might change either because (a) an individual alters his or her absolute preferences, or (b) an individual with fixed preferences encounters changing government policy and thus alters back and forth between more and less to maintain a fixed position in light of a changing government. While empirically separating absolute and relative changes is a challenging task, we would like here to separate them into ideal types to clarify their theoretical standing.

1.2.2 Relative Preference Change

For theoretical clarity, assume that people have fixed (absolute) preferences over policy options. Assume also that government policy changes with changes in party control of the White House. Thus a rational public will change relative attitudes to accommodate changed policy with its position.

How does this micro theory of the loosely informed citizen responding to changing party cues and control of government yield predictions about the shape of aggregate public opinion? That is our task now, turning theory into model.