Partisan Identity in Multi-Party Democracies

Ryan Block

Randy Stevenson

and

Seonghui Lee

Contents

1. Introduction

In this Element, we argue that partisan identification in multiparty democracies is best conceptualized, at the individual-level, as the extent to which a person identifies (positively or negatively) with *each* of the parties in their system. Conceptualizing multi-party partisanship in this way opens new avenues of inquiry into the comparative political behavior of citizens in multi-party democracies and promises new ways to approach several long-standing debates in the literature on partisanship in such systems – e.g., why party identification seems to "move with the vote" in some multi-party democracies or why overall levels of partisanship appear to be declining in some countries. Despite some previous (though little known) work on the topic, the concept of multi-party partisanship is both theoretically and empirically underdeveloped.

One of the goals of this Element is to provide a set of conceptual tools for thinking about multi-party partisanship -- e.g., what is an individual's partisan strength and direction in such a setting? Another is to lay out a theoretical agenda for work on the topic. With this conceptual and theoretical roadmap to guide us, our empirical work will describe, validate, and explore a strategy for efficiently quantifying individual-level multi-party attachments and building measures of concepts like co-partisanship and total partisan strength that are appropriate for multi-party settings.[1] Like much recent work on single-party partisanship, we conceptualize partisanship, or at least that part of partisanship not better captured by other measures, as a social identity – similar (but not identical) to other identities based on class, race, gender, occupation, and nationality. Unlike previous efforts to measure (single-party) partisanship from an identity perspective (e.g., Huddy 2001), our measure can capture multiple, positive or negative, partisan identities of varying strengths or intensities.

To give an immediate sense of how our concept, and its measurement, differs from previous efforts and to provide a preview of the arguments to come, consider our estimate of the partisan identity of a respondent, *i*, to a survey we fielded in Sweden in 2022:

$$\vec{\boldsymbol{\theta}}_{\boldsymbol{i}} = \{2.19, 2.34, 0.73, 1.04, 1.01, -0.41, -0.43, -0.64\},$$

where $\vec{\boldsymbol{\theta}}_{\boldsymbol{i}}$ is a vector and there is one component of this vector, θ_{ij}, for each of the $j \in J$ parties.[2] Below we will argue that larger positive values of θ_{ij} indicate stronger positive identification with party j, while more negative values indicate stronger negative identification with party j. In this conceptualization, a full statement of this respondent's partisan identity is the magnitude and direction of this vector – i.e., a directed arrow in 8-dimensional space stretching from the origin at zero to the point defined by this vector.

---

[1] We use the terms partisan identity and partisan attachment interchangeably. We define $\vec{\boldsymbol{\theta}}_{\boldsymbol{i}}$ more formally in section 3.3.2.

[2] The order of these parties (by average left-right placements) is Left Party, Green Party, Social Democrats, Centre Party, Liberals, Christian Democrats, Moderate Party, and Sweden Democrats.

Figure 1 visualizes this *partisan identity vector* with confidence intervals for each θ_{ij}. In the left panel, we align the parties from left to right based on this respondent's placement of each party on an 11-point left-right scale. In the right panel, the x-axis is based on the average left-right positions of the parties (over the whole sample). The y-axis represents the value of θ_{ij}, our measure of the strength of respondent *i*'s positive or negative attachment to party *j*.

[Figure 1 here]

If this vector captures this person's partisan identity (as we will argue), it paints a clear picture of her affective orientation to the Swedish party system in 2022: She is positively attached to the parties on the left, but less so for moderates compared to more extreme parties. At the same time, she is largely indifferent (if somewhat negative) to the parties on the right, except for a significant negative attachment to the Sweden Democrats – a pattern commonly found in our sample and also noted by Holmberg and Oscarsson (2020) and Bankert (2020) in their separate analyses of negative partisanship in Sweden.

This case also illustrates one of the many drawbacks of applying single-party, "maximalist" measures (and theories) of party identification in a multi-party context.[3] Does this person "think of herself as closer" to the Greens or to the Left Party? Is there truly a party to which she feels closest? If our measure is sound, the answer may well be "no" and thus her response to a maximalist question is unclear: Will she choose the Greens, the Left Party, or simply refuse to make a choice (and so erroneously appear to be independent by the traditional measure).[4]

In the rest of this Element, we explain how we estimate vectors like the one above and argue that they are measuring a specific concept of party identification and so can be used to evaluate theories of party identification that invoke that concept. Further, we show that our measure, while identifying the same partisans as maximalist measures, also reveals a substantial "hidden" partisan affect toward other parties. Crucially, we show that these multiple affective attachments, which vary in strength, are associated with downstream political attitudes and behavior as we would expect if they are measuring the same kinds of attachments (for many parties) that traditional conceptions of party identification measure for a single party: for example, respondents rationalize in favor of the multiple parties to which they are attached, in proportion to the strength of their attachment to each. Finally, we predict and demonstrate that multi-party attachments are particularly likely among parties that form coalition cabinets or otherwise cooperate.

In the rest of this chapter and the next, we outline our theoretical position on the nature of partisan attachments and discuss the conditions that should encourage the development of multi-party attachments. In Chapter 3, we describe the survey battery we have developed to efficiently tap

[3] We define a maximal measure of PID as any question that ask a respondent the party to which they are most attached, closest to, etc.

[4] Even if she does choose one of these parties, will the uncertainty the choice creates result in her reporting a weaker level of attachment than she might report otherwise? As it turns out, in response to the traditional PID question, she chose the Left Party (saying she was "somewhat close" to the Left Party), which was the second highest of her θ_{ij}'s, though her $\theta_{i\,Green}$ and $\theta_{i\,Left\_party}$ are statistically indistinguishable from each other.

respondents' (positive and negative) multi-party attachments and evaluate its practicality for political surveys. We also explain how we apply Item Response Theory (IRT) to transform the responses to our question batteries into estimated attachment vectors (i.e., $\vec{\boldsymbol{\theta}}_i$) for each respondent. Finally, we subject these estimates (and the underlying IRT models) to standard tests of internal scale validity. The results are conclusive: our scale is unidimensional, internally valid, and seems to be measuring the same construct across countries and parties within countries.

Given an internally valid measure of the extent to which our respondents identify positively or negatively with each party, we next explore the structure and usefulness of the measure. Specifically, in Chapter 4, we provide a high-level description of voter types defined by our partisan identity vector $\vec{\boldsymbol{\theta}}_i$. Given our theoretical expectations that multi-party attachments originate from shared policy priorities, patterns of partisan conflict and cooperation, and overlapping social groups (and the role of left-right ideology plays in summarizing all three), we focus on the extent to which our estimates of multi-party attachment are structured by the left-right positions of the parties. As we will see, patterns of multi-partisanship correspond relatively closely, though not entirely, to left-right relationships among parties: The most important multi-partisan identities among our respondents are reasonably described as various kinds of "leftist" or "rightists" (including center-oriented and extremist-oriented versions of both), centrists, and independents. That said, while left-right ideology seems to strongly structure partisan attachments, we also show that they are not equivalent.[5]

Next, in Chapter 5 we ask if our measure of multi-party identification adds anything to traditional conceptions (and measures) of single-party partisan attachment. We begin Section 5.1 with a hard test: Among respondents coded as pure independents by traditional, single-party measures of party identification (no attachment to any party), we examine whether our measure of multiparty attachment is associated with a set of political attitudes and behaviors known to be strongly linked to the strength (and direction) of single-party attachment, vote choice, party likes/dislikes, political participation and engagement, and political efficacy. If our measure of multi-party attachments predicts these political attitudes and behaviors in the expected ways among these traditional independents, we will take that as evidence that our measure taps into the same kind of attachments (for multiple parties) as traditional, single-party, measures do.

In addition to examining traditional independents, Section 5.2 asks whether our measure produces, as maximalist measures typically do, evidence of partisan rationalization or motivated reasoning.[6] We focus on partisan rationalization as the best test of our measure because it has long been considered the "calling card" of partisan identity. That is, voters with an internalized sense of belonging to a partisan group will tend to "adopt the views of the parties and groups they favor" and "develop ideological frameworks rationalizing their group loyalties and denigrating those of their political opponents" (Achen and Bartels 2016: 310). Since we take a social identity

[5] We are not arguing that left-right ideology *causes* partisan attachments – only that left-right can be used to understand the structure those attachments. This structure could easily emerge because other factors cause a respondent to be both similarly attached to two parties and to place them in similar positions on a left-right scale.

[6] We use these terms interchangeably.

approach to conceptualizing partisan attachment, it is important to ask if our measures of multi-party identity produce the same kind of evidence for partisan rationalization (as shown in the literature using traditional measures) and whether the extent of such rationalization is proportional to identity strength for each party across the full range of possible identity strengths. Further, we pay close attention to whether attachments to parties other than the one to which a respondent is *most* attached also produce evidence of partisan rationalization (in proportion to levels of attachment). If they do not, it would raise doubts that our measures of such "secondary" attachments are capturing the same construct as the primary (or traditional) attachment.[7]

With that roadmap, the rest of this chapter asks why we need a multi-dimensional concept (and measure) of partisan attachment. Is it likely to help us settle open theoretical questions or open new avenues of inquiry?

1.1 Why multi-party Identity?

From the beginning, scholars developing the concept of party identification asked whether and how it might apply to multi-party systems (Campbell et al. 1960; Weisberg 1980). Indeed, the first instinct of many European political scientists was that party identification in multi-party systems *must be* multi-valued. For example, following a long tradition of European sociologists like Lipset and Rokkan (1967), van der Eijk and Niemöller (1983) linked overlapping partisan attachments in the Netherlands to the Dutch cleavage structure, demonstrating that about half of all party identifiers had multiple attachments. Likewise, U.S. scholars who think of partisanship as a social identity have often called for multi-party measures. As Schickler and Green (1997) explain:

> "…we cannot assess directly the effectiveness of our approach until the SRC-style party ID question is augmented by questions tailored to multi-party environments. In particular, this means asking about levels of identification one party at a time, much in the same way that party affect is currently measured by a succession of thermometer items. ***It is our hope that survey researchers—whether or not they are persuaded by our thesis concerning partisan stability—will begin to introduce party-by-party measures of identification into the panel studies they conduct***." (Schickler and Green 1997: 478, emphasis added)

Other, more recent, calls for multi-party measures of party identification include Bankert (2020, 2024), Carius-Munz (2020), Mayer and Schultze (2019), Caruana, McGregor, and Stephenson (2015), Garry (2007), Johnston (2006), Oscarsson and Holmberg (2020), Oshri, Yair, and Huddy (2022), Rose and Mishler (1998), Rosema and Mayer (2020), and Schmitt (2009). Despite these calls for theoretical and empirical development of the concept of multi-party identification, progress has been slow. This may be partly due to the influence of some empirical work rejecting

[7] Depending on the context, we use the term "secondary attachments" to mean either all the partisan attachments that are weaker than the respondent's largest ("primary") one, or attachments to parties other than the one identified by traditional measures.

the idea (see the discussion of Schmitt 2009 below), but it is more likely due to the persistent lack of compelling measures of multi-party identification (Garry 2007) – a lack we aim to rectify.

Why do these scholars think we should measure (and understand the nature of) multi-party identifications? The first reason begins with the simple fact that many political attitudes and behaviors work differently in multi-party democracies than in two-party systems. Voters in multi-party democracies in Europe, compared to those in two-party democracies, know more about politics (Milner 2002; Lin, Santoso, and Stevenson 2024) and the ideological positions of parties (Gordon and Segura 1997; Fortunato, Stevenson, and Vonnahme 2016), are more interested in politics (Lee, Stevenson, and Cozza 2025), interpret the left-right ideological dimension differently (Lee, Santoso, and Stevenson 2025), engage in performance voting at different rates and for different reasons (Duch and Stevenson 2008; van der Brug, van der Eijk, and Franklin 2007; Lewis-Beck 1990), access policy-making responsibility differently (Duch, Przepiorka, and Stevenson 2015; Fortunato et al. 2021), and engage in different kinds and amounts of strategic voting (Cox 2018; Gschwend and Meffert 2017). More generally, given the myriad ways mass political behavior differs between two-party and multi-party systems, perhaps the default assumption should always been that partisanship would operate differently in these two distinct political environments.

Indeed, some political behaviors essential to understanding the peculiarities of political behavior in multi-party systems likely depend heavily on voters' multi-party identities. For example, it is hard to imagine that voters' preferences over potential post-election governing coalitions (and their retrospective evaluations of incumbent coalitions) do not depend, at least in part, on how strongly they identify (positively or negatively) with the various coalition partners (Hahm, Hilpert, and König 2024). A similar argument applies to studies of strategic, coalition-directed voting. Further, although an extensive literature extolling the virtues of proportional, multi-party coalition government for promoting compromise among parties and their partisans (e.g., Lijphart 2012), scholars have only very recently begun to explore whether one of the palliative supporting such compromise is the development of co-partisanship among supporters of allied parties (Hahm, Hilpert, and König 2024; Horne, Adams, and Gidron 2023; Gidron, Adams, and Horne 2023).

A second reason to study multi-party identification is that the small number of empirical studies that have already tried to measure multi-party attachments have found them to be common, consequential, and growing over time. To our knowledge, the first such empirical study was van der Eijk and Niemöller's (1983) study of Dutch voters, which explicitly asked survey respondents if they were multiply attached to parties (without a multi-item attachment battery) and found widespread multi-party attachments. Much later, Garry (2007) also demonstrated, using questions about each Northern Irish party, that 31% of positive party identifiers had positive attachments to more than one party on the same side of the Protestant/Catholic divide. Further, 57% of negative identifiers held negative attachments to multiple parties on the opposing side.

Likewise, even though Schmitt (2009) rejected van der Eijk and Niemöller's (1983) finding that half of all Dutch identifiers held multiple partisan identities – because few respondents to the Dutch CSES surveys *spontaneously* indicated multiple PIDs when unprompted – his own results for several other countries (even under this *very* strict standard) showed high levels of spontaneously volunteered multi-party attachments. For example, Schmitt reports that 22%, 16%, and 10% of Norwegian, Swedish, and German respondents to the first CSES module (1996-2000) spontaneously reported multi-party attachments. We thus agree with Garry (2007) that Schmitt's

analysis actually bolsters the case for the existence of multi-party attachments. Finally, other scholars, lacking explicit measures of multi-party attachments, have used party like/dislike scores to proxy for multi-party attachments. For instance, Holmberg and Oscarsson (2020) use Swedish National Election Study data to show that from 1979 to 2018, there was a dramatic increase (by nearly 40%) in the number of Swedes that assign the same like/dislike rating to two or more parties – a finding they interpret as consistent with the idea that multi-party identification may be growing over time.

A third reason to study (and measure) multi-party attachments is that they suggest an alternative theoretical explanation for the frequent claim in the literature that partisan strength or intensity (using single-party concepts) has waned over time (either in specific countries or cross-nationally). For example, Holmberg and Oscarsson (2020: 19) use data on 45,000 Swedes, surveyed from 1968 to 2018, to conclude that Swedish voters "have, during the last 50 years, tended to be become less identified with the parties they vote for." Lacking a robust theoretical framework for multi-party identity, however, they never considered an otherwise obvious alternative explanation for this trend: the strength of partisan identification may not be declining as much as it is being "redistributed," within individuals, across a wider set of parties. If, for example, the average adherent of the Social Democrats in 1968 has become less strongly attached to that party *because* she has redistributed some of her partisan attachment to Greens or Left Party, her overall level of partisan attachment could remain stable (though now spread-out over more parties), even as her single-party attachment to the Social Democrats declines.[8] If this pattern described many individuals who were once single-party partisans (of various parties), the total number of people reporting a strong attachment to a single party (in a single-party, maximalist question) would certainly decline.[9] Use of a multi-item attachment battery, asked about each party in the system, is likely the only way to adjudicate between these possible explanations.[10]

Fourth, incorporating multiple partisan identities into the theoretical arsenal of comparative political behavior promises to help resolve the long-standing critique that, in some multi-party democracies, partisan identity is not stable overtime and/or that it follows the vote (e.g., Thomassen 1976). While this claim is often over-generalized (e.g., see the helpful discussion in Oscarsson and Holmberg 2020), its influence on the field of comparative political behavior is undeniable. Indeed, doubts about the relevance of party identification in multi-party democracies are likely one reason multi-party identities have for so long been (mostly) ignored.

In our view, however, it is the very possibility of multi-party identities that may resolve the long-standing tension between the outsized role of party identification in explaining the political behavior of Americans and its more limited role in explaining political behavior in (some) multi-

[8] We do not mean to imply that attachment is a limited quantity for each individual. However, whatever process governs multiple attachments, it would not be surprising if it includes mechanisms that tend to dilute the strength of partisan feelings for one party as attachment to others strengthens.
[9] It's also likely that some multiply attached respondents who are forced to answer a single-party maximalist PID question will refuse to do so and thus falsely appear to be independents.
[10] The same process may also account for the finding that people are more likely to form a partisan bond in systems with fewer parties (Berglund et al. 2005; Önnudóttir and Harðarson 2020).

party democracies. Specifically, once we reconceive party identification as a vector valued quantity, simple theoretical models can show how voters with multi-party identities use them in vote choice, modify them due to short-term electoral factors, and importantly, report them when asked a maximalist PID question. These models predict exactly the fluidity often observed in reported partisan identification in multi-party democracies, even if the underlying vector of identities remains unchanged. That said, since the, exclusively cross-sectional, empirical work in this Element is not well-suited to testing the inherently dynamic hypotheses that come out of such models (though see our concluding section for one limited attempt to do so), we relegate the exposition of that model to Appendix E and a more thorough discussion of its implications and potential critical tests to our discussion, in the conclusion, of future work.[11]

Finally, a healthy literature on multi-party attachments is likely to lead to new questions and avenues of inquiry that are not available otherwise. For example, what is the relationship between multi-party identities and the kind of hierarchical, ideological identities (e.g., "leftist") that scholars have begun to explore (González et al. 2008, Hagevi 2015, Lin, Santoso, and Stevenson 2025)? Likewise, if multiple partisan identities exist, how does an individual manage them and adjudicate between them when they are in conflict? Social psychologists have a rich literature asking such questions for other kinds of social identities (see the review by Kulich et al. 2017), but, lacking a significant literature on multi-party identification, this work has not yet penetrated political science. While some political scientists have certainly contemplated multi-party identity, no work has explored the theoretical and conceptual subtleties that arise in using the concept. For example, what is the best way to conceptualize co-partisanship or an individual's overall partisan strength when that strength may be distributed across multiple parties?

[11] To preview, the main implication of the theory is that the observed temporal instability of reported, maximalist measures of single-party identification (in places like the Netherlands) should be greatly reduced if one instead measures the underlying vector $\vec{\boldsymbol{\theta}}_{\boldsymbol{i}}$ over time. A proper test of this will thus require the collection of (expensive) panel data. In this Element, we hope to justify future investments into collecting such data by addressing some of the key theoretical, conceptual, and measurement questions that arise when one takes the idea seriously.

2. The nature and origins of partisan identification

In this chapter, we explain the theoretical understanding of partisan identification that guided the development of our proposed measure and shaped our expectations about what downstream political beliefs and behaviors it should help explain. Once we are clear about the concept of partisan identification, we can then productively ask how to consistently conceptualize and measure other related concepts like negative partisanship, independence, overall partisan strength, co-partisanship, and multi-party identification.

Several recent reviews detail the intellectual history and conceptual development of party identification (e.g., Jost, Baldassarri, and Druckman 2022; Oscarsson and Holmberg 2020, Groenendyk 2013; Lavine, Johnston, & Steenbergen 2012; Sniderman and Stiglitz 2012). Almost all work defines the concept in one of two ways. One camp sees partisan identification as a social identity. This was the original concept of party identification that animated the *American Voter* (1960) and has since been updated using the theoretical apparatus of social identity theory (Tajfel 1981), by scholars like Green, Palmquist, and Schickler (2002), Greene (1999, 2004), Huddy and Bankert (2017), Mason (2015), Rosema and Mayer (2020). A second, "revisionist" tradition conceptualizes partisan identification as a running evaluation of a party's policies and performance (e.g., Fiorina 1981, Achen 1992).[12] Finally, theoretical attempts to reconcile the two approaches – including theories of partisan ambivalence (Lavine et al. 2012), party reputations (Sniderman and Stiglitz 2012), dual or triple motivation theories (Groenendyk 2013, Jost, Baldassarri, and Druckman 2022), as well as reviews of the relevant empirical evidence (e.g., Shi et al. 2024) – have concluded that "the evidence shows that both views of partisanship are correct [...] partisanship functions both as a stable psychological construct involving affective group attachment and as a temporary political summary judgment of group performance." (Levendusky 2017: 90, quoting in part from Lavine, Johnson, and Steenbergen 2012:10).

Below, we examine these two traditions, and with that background, explain our own theoretical take on the nature of partisanship and how that informs our thinking about multi-party partisanship. Before doing so, however, we propose and adopt an innovation in nomenclature to maintain conceptual clarity that is sometimes lacking in work on party identification: we separate the idea of party identification as *partisan social identity* from the concept of party *policy and performance evaluations,* labelling them "PSI" and "PPE," respectively.[13] These labels emphasize the key distinction: *PSI is a social identity,* and *PPE is a bundle of party evaluations*. In what follows, we use the generic label "PID" when we, or the works we are discussing, do not distinguish between these different conceptions of partisan identification. However, when our (or other's) theorizing or empirical work is focused clearly on the concept of party identification as a social identity, we will refer it as "PSI," and when we discuss partisan identification as a bundle of performance and policy evaluations, we use "PPE." We should not, however, take this distinction too far. As we explain in

[12] We think of performance broadly to include leader competence and the like. Similarly, we use the term policy to mean party positions - relative to the voter- on all manner of narrow issues as well as core values.
[13] See Appendix F.1.

Section 2.2, even in the relatively pristine environment of the lab, researchers have struggled to manipulate PSI and PPE separately, as citizens' evaluations on party policy and performance appear inexorably linked to their expressions of identity (Mason 2018) and vice versa (Orr, Fowler, and Huber 2023).

Consequently, we want to make clear that, at this stage, we are not trying to settle (or even empirically contribute to) the on-going debate about whether identity or policy/performance evaluations are more central to party identification. Instead, our goal for this Element is simply to explore the usefulness of multi-party identification as a concept and build a measure of it that serves as the multi-party equivalent to the measures of single-party identification commonly used by scholars working from either perspective. As we argue below, that does not require us to be equivocal about the concept we are trying to measure or the instrumentation that most effectively measures it. If partisanship has two faces – PSI and PPE – then scholars should, in our view, try to measure both (as we do below), while recognizing that PSI and PPE may be less distinct in the minds of voters than scholars, seeking clear conceptual categories, might hope.

Indeed, multi-party measures of voters' *evaluations* of parties' policies and performance are already included in most election surveys. Likewise, scholars have explored the nature of those evaluations and their impact on voter behavior in multi-party – e.g., on attributions of responsibility across coalition partners (Bowler, Gschwend, and Indridason 2018; Fortunato et al. 2020) and vote choice (Duch and Stevenson 2008; van der Brug, van der Eijk, and Franklin 2007). In contrast, the literature includes few discussions of multi-party partisan identity, including its development, stability over time or over the life-cycle, its measurement, or its consequences. Given this, we argue that what is needed is more theory and instrumentation aimed at understanding and measuring multi-party *PSI*, which can be used in combination with various measures of PPE in a holistic assessment of multi-party identification. Ultimately, we hope this Element will begin a conversation about the usefulness of multi-party partisanship as a concept and how best to measure it, thereby contributing to the ongoing debate about the nature of partisanship more broadly.

2.1 Partisanship as a social identity

A great deal of research conceptualizes party identification as a social identity. A social identity involves a subjective sense of group belonging, internalized to varying degrees, which results in individual differences in identity strength, a desire to positively distinguish the group from others, and the development of in-group bias (e.g., Huddy, Bankert, and Davies 2018; Tajfel 1981). Like other social identities, the key consequence of this conceptualization is the individual's self-conception as a member of a partisan in-group. The psychological consequences of such identities are well-documented by social psychologists. Individuals who identify strongly with a social group are likely to derive emotional benefits from perceived increases in group status (Ellemers, Spears, and Doosje 1997), work to maintain positive stereotypes about the group (Andreychik and Gill 2009), derogate relevant out-groups (see Branscombe et al. 1999), and hold biased beliefs about variables relevant to group status. These differences between in and out group members have been substantiated repeatedly in lab and field studies (e.g., Allamong and Peterson 2021; Iyengar and

Krupenkin 2018; Linville, Fischer, and Salovey 1989; Rothschild et al. 2019; Lane 2015) and, in some cases, have been traced to changes in patterns of brain activation (Hein et al. 2010, and especially Harris et al., 2022).

Building on the large social psychological literature on social identity, scholars have generated a great deal of evidence that partisan identification is usefully understood as a social identity and that its unique content and impact arises *because* it is a social identification (Green, Palmquist, and Schickler 2002; Greene 1999, 2002, 2004; Weisberg and Greene 2003; Lavine, Johnston, and Steenbergen 2012; Groenendyk 2013; Huddy, Mason, and Aarøe 2015; Achen and Bartels 2016; Huddy, Bankert, and Davies 2018).

Despite some recent doubts (e.g., Fowler 2020; Orr, Fowler, and Huber 2023), the continuing importance of party identification to modern debates about political behavior – and its roots in identity politics – is reflected in two major critiques of the rationality and knowledge of American voters. Specifically, both Achen and Bartels (2016) and Kinder and Kalmoe (2017) argue that most Americans (~70%) lack coherent or stable issue preferences and do not reliably link policies to any sort of ideological description of parties (e.g., liberal-conservative).[14] Further, Achen and Bartels challenge the idea that most voters can reliably use party performance evaluations to assess party records in sensible ways.

Given that ~70% of citizens lack coherent policy and performance evaluations, how can we explain their political behavior? Both Kinder and Kalmoe and Achen and Bartels point to group-based models in which various group identities are central, and partisanship is one of the most important identities driving political behavior. Achen and Bartels (p. 299) succinctly conclude: "For most citizens most of the time, party and group loyalties are the primary drivers of vote choices." Further, these authors and many others argue that when less sophisticated partisans espouse sensible policy positions, they are often simply following cues provided by partisan elites from the parties to which they are attached (e.g., Meyer and Wagner 2020, Duch and Stevenson 2010).

A dated body of work in comparative politics expresses similar worries about citizen sophistication in other western democracies (e.g., Butler and Stokes 1969; Converse and Pierce 1986; Fuchs and Klingemann 1990; Inglehart 1990; Klingemann 1979; Putnam, Leonardi, and Nanetti 1979). More recent work, however, has softened this critique for voters in multi-party systems, suggesting that the necessity of governing coalitions and greater opportunities for coalition-directed strategic voting in these systems provide stronger incentives for citizens to learn specific kinds of political information that are less relevant in the American system (Fortunato and Stevenson 2013): Voters in coalition systems seem to have high levels of knowledge of the broad left-right relationships between of parties and their relative sizes (Fortunato, Stevenson, and Vonnahme 2016; Haime, Lee,

[14] See Appendix F.2 for discussion of Jost's (2021) disagreement with Kinder and Kalmoe (and Achen and Bartels) about the importance of ideology generally and about the comparative work on political knowledge that questions the universality of the "know-nothing" voter.

and Stevenson 2017).[15] Even so, however, this work does not claim most voters have detailed information about party policies or performance. Instead, it argues that low-information voters can leverage relatively cheap information cues, like parties' relative left-right images, to make sensible inferences about more complicated aspects of coalition politics, such as attributing responsibility to the parties in coalition (Fortunato et al. 2021) or predicting which parties will form a coalition after elections (Fortunato and Stevenson 2025).

Finally, many researchers adopting the social identity approach posit that partisans think and act differently from non-partisans because they have a set of experiences that have quite literally changed how their brains function (see the useful reviews by Bavel and Pereira 2018 and Harris et al. 2022). While neuroscientific studies of partisanship remain rare, several support the hypothesis that partisan identity is substantiated in the brain other social identities (Cikara and Bavel 2014; Cikara et al. 2017) and even "brand" loyalties (McClure et al. 2004), by altering how partisans process information compared to non-partisans (Hein et al. 2010; Hibbing, Smith, and Alford 2013).

2.2 Partisanship as policy and performance evaluations

Franklin and Jackson (1983: 958) nicely summarize the "revisionist" theory of party identification put forward, most famously, by Fiorina (1978, 1981) and adopted by a great deal of subsequent scholarship:

> *"Party preference was treated in this model as a function of voters' issue positions and the positions of the parties. The cost of information might lead voters to adopt methods of evaluating party positions which could lag behind actual party policies, and this could produce apparent loyalty to a party. However, once provided with more current information, voters would immediately reevaluate their party preferences. Thus, there was no party loyalty per se in the Downsian framework. Stability of party preferences would result only from stable policy preferences and party platforms."*

This quote captures how some revisionists, in providing an alternative theory of party identification, actually denied its overall relevance as a distinct concept. If party identification is nothing but policy and performance evaluations, why would we need a separate concept of PID? Surely, we should simply theorize about and measure the relevant evaluations and include them in our empirical models of vote choice and other behaviors without the need for party identification as a separate concept.[16]

Empirically, the revisionist argument has long relied on observational work challenging the notion of partisanship as the "unmoved mover" that Campbell et al. (1960) claimed. These studies demonstrate that the strength (though less often direction) of partisan identity responds to changes in policy preferences (Carsey and Layman 2006; Dancey and Goren 2010; Franklin 1992; Franklin

[15] For example, Fortunato, Stevenson, and Vonnahme (2016: 1211) report results from a Danish national election survey in which more than half the respondents placed all 15 pairs of Danish parties in the correct left-right order.

[16] Indeed, the huge literatures on the spatial model of voting and performance voting do exactly that.

and Jackson 1983; Jackson 1975; Highton and Kam 2011), performance evaluations (Fiorina 1981; Weinschenk 2010), candidate evaluations (Page and Jones 1979), and past votes (Markus and Converse 1979). Revisionists scholars also showed that partisanship varies over time at the individual level (e.g., Brody and Rothenberg 1988).[17]

2.3. Where does this leave us?

The above review makes clear that important debates about the nature of partisanship remain unsolved, and we do not aim to resolve them here. What, then, should we take from them?[18] First, much of this debate is empirical and a fair reading of the empirical work reveals that voters' partisan attachments function both as a social identity and a running evaluation of parties' policies and performance. From a theoretical perspective, however, it seems clear that partisanship is only interesting as a concept, separate from party policy and performance evaluations, if it is, at least in part, a social identity. If it were not, there would be no reason to distinguish how a partisan makes performance evaluations from how an equally-informed non-partisan does the same.

Most political scientists – even if skeptical of the overarching importance of partisanship – would reject this view. Instead, most would grant that identifying with a partisan group changes an individual's relationship to the party system relative to a non-partisan, by making them more sensitive to partisan information, altering stereotypical beliefs about partisan identified objects, changing perceptions of her own relationship to the parties, biasing information processing, and more. The importance of these impacts may vary across individuals. Ultimately, however, it is partisanship as an identity that animates and gives force to partisanship as a concept separate from party evaluations. Thus, in our view, it is partisanship *as an identity* that we want to study and measure.

That said, we do not ignore the lessons of the partisanship-as-PPE literature. Most importantly, we do not pretend our measure of PSI is "untainted" by any evaluations. For example, we accept that answers to party identification questions likely include an implicit evaluation of one's ideological relationship with the party. Indeed, as discussed below, given the programmatic nature of (much of) party politics (at least in Western democracies), party identification often develops in the context of (at least some) policy concerns. Likewise, once formed, this identity may lead individuals to adopt party cues about other policies and performance evaluations, bundling these concerns up with other aspects of identity into their response to a partisan identity item.

Having adopted this identity perspective, in the next section we ask how such identities develop, and most importantly, how co-partisan or multi-partisan identities can develop. This will lead to hypotheses about both the *kinds of parties* that should have the most co-partisans and the *contexts* in which we should see more or less co-partisanship.

---

[17] See Appendix F.3 for a more thorough discussion.
[18] See Appendix F.4.

2.4 The development and structure of partisan identity in multi-party democracies

Two of our goals in this Element are to introduce our measure of multi-party partisan identity and to use it to explore the structure of these identities (e.g., to which pairs of parties would we expect voters to develop co-partisan attachments?). To do that, in this section we discuss why individuals might develop and maintain such identities. Specifically, we discuss three pathways that can lead to multi-party PSI and that should structure multi-party identities across individuals. The first two are familiar, because they are straightforward extensions of the same factors explaining the development (and maintenance) of single-party identity: the match between one's own values and policy preferences and those of the parties, and the match between one's other social identities and those of the parties' leaders and supporters. The third factor is more unique to the development of party identification in multi-party systems and may be particularly important for multiple partisan identities: the ongoing need for parties to cooperate in governing and/or policy-making coalitions. Finally, after exploring these three drivers of multiple partisan identities, we examine the central role that left-right ideology plays in aggregating these factors in the minds of voters.

2.4.1 What we know about the factors that create and maintain single-party partisan identities

Writing in 1996, Miller and Shanks opined, "Despite the profound importance of party identification [...], the precise origins of individual party loyalties are not well understood." (p.128) Unfortunately, the situation has not much improved since. Early work (e.g., Sears and Funk 1999; Converse 1964; Campbell et al. 1960) characterized PID as developing within the family, in part due to the strong empirical association between individual partisanship and *recalled* parental partisanship (by the adults being surveyed). However, stronger research designs that examined parents and offspring at the same time and/or examined other influences outside the home revealed a much weaker connection between parental socialization and partisan attachments in the U.S. and a correspondingly greater influence of peers, teachers, and other life experiences (e.g. Jennings and Niemi 1975; Niemi and Jennings 1991; Niemi, Katz, and Newman 1980).

Further, these identities, while not frozen over the life-course, are sticky for many voters (Clarke and McCutcheon 2009; Green and Palmquist 1994). Thus, while PID changes in response to contemporaneous events, for example, due to updated policy and performance evaluations (Fiorina 1981; Neimi and Weisberg 2001), most scholars concluded these effects are smaller than the reverse and are highly circumscribed. As Dassonneville, Fournier, and Somer-Topcu (2023) put it in a recent review:

> *"...most switching takes place between identifying with one specific party and independence (Neundorf et al. 2011; Zuckerman and Kroh 2006). Changes, furthermore, are*

> *neither rapid nor universal,' which has led Tucker et al. (2019: 324) to characterize partisanship as a 'very slow mover'. Changes in party identification thus do occur, even if partisan attachments are not fickle."*

Oscarsson and Holmberg's (2020: 19) review comes to similar conclusions:

> *"... data tend to lend support to [a model] which states that people's sentiments are primarily formed at an early age and remain fairly stable thereafter. Studies performed in America (Abramson 1983), the UK (Crewe and Thomson 1999) and Sweden (Holmberg and Oscarsson 2004; Westholm 1991) indicate that voters' levels of party identification remain stable... over time."*

If we think of partisanship as a social identity, then one answer to how it is maintained or sustained (i.e., why it is sticky) is immediate: changes to the way partisans think and perceive the world produce a sequence of (partially biased) inferences about the world that reinforce the identity (Campbell 1960; Stokes 1966: 127; Lavine, Johnston, and Steenbergen 2013). These identities can certainly be shocked (e.g., realignment events) and respond to short-term influences like leader or performance evaluations (which are themselves likely biased by PSI). However, in the normal course of events, the psychological process of identity maintenance is likely enough to produce the level of continuity in partisan identities we observe. Indeed, once we allow for multiple identities, this finding based on the U.S-case can be extended to multi-party systems, in which the temporal stability of PID has often been doubted.

If the psychological changes that constitute PSI explain (some amount of) stability in partisan attachments, this leaves the question of how they develop initially – which, again, matters because the factors that create multiple party identities are likely the very ones that will help us to discover how – if at all – they are structured.

2.4.2 Shared policy goals

The idea that an individual may be attracted to a party because that party advocates for policies with which she agrees is surprisingly rare in the scholarly literature on the subject, perhaps because the idea plays the role of foil in many discussions. Indeed, the parental socialization literature discusses the development of PID as if it is transmitted from parent to child free of policy content. Likewise, the authors of *The American Voter* posit that policy views were developed after a party identity was acquired – and that largely determined those views.

The work that most consistently recognizes a role for policy in the development of party identification actually comes from revisionists (e.g., Franklin and Jackson 1983; Fiorina 1981; Carsey and Layman 2006), who largely denied the conception of PID as social identity. As a result of this, one can hardly lean on the PID-as-PPE literature to argue that people develop a partisan social identity (PSI) because they agree with party's policy agenda. Instead, there remains a long-standing, implicit understanding that policy probably matters to the development of PID, but that

such explanations can only succeed if they answer the question of how the typical voter comes to know enough about policy debates (when they are not yet a partisan).[19]

As such, some scholars argue that for PID, it is not many narrow policies, but broad sets of core values that matter. A typical voter may be unaware of parties' specific policy positions on many different narrow policies; but, they know, for example, that the Socialists are for equality and the Conservatives less so. One important body of work invoking such argument is the comparative literature on party brands (e.g., Baker et al. 2016; Lupu 2016). As Evans and Neundorf's (2020: 1264) summarize:

> *"If people hold fundamental and enduring attitudes towards economic and political principles, such as equality, that influence their attitudes towards political issues, it is also likely that these values can shape their partisanship. Such central elements of political belief systems can be expected to influence party preferences as voters update their partisan identities to correspond with their values: if values 'predispose us to favor one particular political or religious ideology over another' (Rokeach 1973: 13), it is plausible that they can predispose people to favor one political party over another."*

Thus, for the ~30% of citizens who are inherently interested and knowledgeable about politics (Kinder and Kalmoe 2017; Prior 2019) as well as for those who, while less knowledgeable, grasp the broad value differences between parties, the development of partisan identities may well depend on the match between voters' and parties' broad values.

Applied to multi-party identification, this has an obvious implication: holding other factors constant, we should expect individuals to develop similar attachments to parties with similar policy agendas and/or broad values. Given the greater similarity among parties in their core values compared to their specific policy commitments, this association may be most consequential for the latter.

This resonates with recent work in social psychology (Roth et al. 2025) that has tried to identify multiple identities that an individual is likely to develop and maintain, arguing that two groups to which a person maintains equally strong social identities must have overlapping stereotypical traits (e.g., in our case, shared policy goals). Specifically, the Multiple Identity Integration Model (MIIM; Roth et al. 2025) is an extension to multiple identities of Greenwald et al.'s (2002) model of "balanced identity theory." That theory argues that any identity with a social group is strengthened when the individual's self-concept includes (positive) traits that are also stereotypical characteristics of a social group. Extended to multiple groups (and invoking the principle of cognitive consistency), the MIIM implies that a person's joint attachment to two groups is strengthened when the two groups share stereotypical group traits that are also part of the person's self-concept (Irish people are religious; Catholics are religious; I am religious; so, my identification with Catholics and the Irish is strengthened). Applied to party attachments, the MIIM model would

[19] See Appendix F.5.

thus predict that if two parties have a great deal of policy overlap (a stereotypical trait) then identities to both parties should be strengthened.

2.4.3 Overlapping social identities

Another prominent view of the development of partisan attachments relies on the observation that political parties tend to consist of people from the same (often overlapping) set of social groups. Political scientists have suggested that one path to a specific partisan identity is a kind of matching process in which an individual compares her own portfolio of social identities to the mix of social groups that tend to identify with the party (e.g., Green, Palmquist, and Schickler 2002; Mason 2018; Huddy 2001; Harteveld 2021). Green, Palmquist, and Schickler (2002: 7) explain it this way:

> *"Social identification involves comparing a judgment about oneself with one's perceptions of a social group. As people reflect on whether they are Democrats or Republicans (or neither), they call to mind some mental image, or stereotype, of what these sorts of people are like and square those images with their own self-conceptions. In effect, people ask themselves two questions: What kinds of social groups come to mind as I think about Democrats and Republicans, and Independents? Which assemblage of groups (if any) best describes me?"*

Applied to the development of multi-party identity, this implies that a person should develop co-partisan attachments between pairs of parties who share the same mix of supporting social groups. Likewise, a person should develop negative attachments to parties mainly composed of her relevant "out-groups." [20, 21]

2.4.4 Patterns of partisan cooperation and conflict

Multi-party systems in which coalition cabinets often form (or where less formal multi-party legislative coalitions are necessary to govern) naturally encourage citizens to think in terms of partisan allies and enemies. In many cases, elections in these systems are fought explicitly in terms of such coalitions – even conveniently labelling them for voters (e.g., Italy's "Olive Coalition" and Sweden's "Red-Greens" on the left and "The Alliance" on the right). Thus, we argue that coalition politics provides a ready pathway for voters to develop attachments across party lines.

While there has not been a lot of work that explores how the experience of coalition government affects partisans' feelings toward their partners, recent studies have begun to do so.[22] For example, Horne, Adams, and Gidron (2023) examined the empirical association between formal cooperation

[20] See Appendix F.6.
[21] See Appendix F.7.
[22] Gaertner, et al. (1989) is an early social psychological study of group identity in which they argue for (and find in a laboratory experiment) that the usual biases associated with in-group membership are ameliorated when subjects are manipulated to think of their group as part of a larger identity group.

in cabinet and the like/dislike ratings partisans give to their coalition partners across 19 Western democracies between 1996 and 2017, concluding that:

> *"…elite inter-party cooperation in consensual systems is associated with reduced inter-party hostility in the mass public. This is because governing parties' supporters feel much more warmly toward their coalition partner(s) than we can explain based on policy agreement alone. Moreover, these warm affective evaluations linger long after the coalition itself has dissolved."*

Likewise, Bassan-Nygate and Weiss (2021), leveraging post-election uncertainty about cabinet formation following the 2019 Israeli legislative election, conducted a survey experiment that manipulated respondents' perceptions of the likelihood that a unity government would form, finding that priming party cooperation promoted inter-party tolerance.

More recently, Hahm, Hilbert, and König (2023) used a conjoint experiment, randomly assigning characteristics (age, gender, religion, nationality, social class, and partisanship) to a fictitious partner in a set of economic games. They found that partisanship influenced allocations more strongly than other characteristics, and that this effect was conditional on the history of coalition partnership between the relevant parties.

Other work suggests that frequent party cooperation may lead voters to develop superordinate identities in which they come to identify with the whole bloc of parties as a group (Lin, Santoso, and Stevenson 2025; Hagevi 2015). Hagevi (2015: 74), for example, argues that efforts to build political blocs "could result in emotional and intellectual attachments to these blocs in the electorate." Further, trends toward more bloc electoral competition in many European counties, combined with powerful modern campaign communications and the "presidentialization" of electoral competition, likely encourages such cross-party attachments.[23]

This argument leads to two implications that we can pursue with a measure of multi-party PSI. First, to the extent partisan cooperation and conflict drives multi-party attachments, such attachments should be rarer (at least on the positive side) in democracies like the UK, where governing coalitions are uncommon. We examine this hypothesis empirically in Chapter 4. Second, if co-partisanship for two parties arises (at least in part) because those parties have a history of cooperation, we should find that an appropriate measure of co-partisanship, which we develop in section 3.2.5.3 , will covary with parties' histories of cooperation. Using our estimated vectors for multi-party attachments, in section 4.3 we analyze this relationship for Denmark and the UK, using a survey question that asks respondents the extent to which each pair of parties in their system cooperate with each other (Lee, Santoso, and Stevenson, 2025).

[23] See Appendix F.8.

2.4.5 Left-Right ideology and the structure of multi-party identification

We have argued that multi-party attachments arise because pairs of parties have overlapping values, policy agendas, and social group support. Further, many pairs of parties have long and well-established patterns of cooperation (or conflict), which facilitate (or impede) co-partisanship among their supporters. Recent work exploring how citizens in multi-party democracies think about the relative left-right positions of parties has shown that citizens effectively and efficiently summarize these three influences in their placements of parties on the left-right (Lee, Santoso, and Stevenson 2025).[24] This suggests that we can use the left-right ordering of (or distances between) parties as a summary of the four forces driving co-partisanship, and thus expect the left-right to strongly structure multi-party attachments.[25]

Finally, it is important to emphasize that we are not making a strong claim about the causal ordering of changes in PSI, left-right placements, and beliefs about partisan conflict and cooperation, partisan policy, or social group support. After all, the purpose of studying PSI as a distinct concept is that changes in the strength or direction of partisan identity affect how a person thinks about all these other things, and vice versa.[26, 27]

2.5 Negative partisan identity

In recent years, scholars, politicians, and the media have become fascinated with the idea of negative PID (Abramowitz and Webster 2016; Caruana, McGregor, and Stephenson 2015; Iyengar, Sood, and Lelkes 2012; Mason 2015; McGregor, Caruana, and Stephenson 2015; Medeiros and Noël 2014). While the concept has been a part of the conversation about partisan identification since the beginning (e.g., Campbell et al. 1960, p. 121; Campbell, Gurin, and Miller 1954), it has seldom been carefully defined (Rosema and Mayer 2020:127).[28] Usefully, however, Bankert (2024) provides an extensive discussion of negative PID as a social identity, arguing that a negative social identity is essentially the same as a positive one, but "centers on the rejection of a group and its members." Theoretically, the question is whether people really think of themselves as part of a group opposed to a given party and whether they feel *connected* to its members? Would they subconsciously favor members of this negatively defined group? Would they be more likely to believe messages consistent with the stereotypical world view of members of such a group? To Bankert and many others, there are clear examples in which the answers to these questions are *yes* – e.g., many Swedes appear to identify with the group of people who oppose the Sweden

[24] See Appendix F.9.
[25] See Appendix F.10, F11, and F12.
[26] See Appendix F.13.
[27] The inclusion of "direction" here makes a lot more sense when we allow multi-party identification, since it's easy to imagine the direction of this vector (and its length) changing much more subtly over time than the traditional conception in which changes in partisan direction require large, discrete jumps between parties.
[28] See Appendix F.14.

Democrats (Bankert 2024) and many Americans with those who oppose MAGA Republicans, even if they do not strongly identify with other parties.[29]

Further, in her analysis of a multi-item negative attachment scale administered in the US, UK, Italy, the Netherlands, and Sweden, Bankert's respondents often revealed negative attachments to parties that were not just reflections of positive PSI – they existed "without equally strong attachments to an in-party" (p.35). Likewise, earlier work in Europe, the U.S., and Canada used party "like/dislike" scales or questions about whether a respondent "would never vote for" a party to draw similar conclusions about the potential non-dependence of in-party and out-party affect (Caruana, McGregor, and Stephenson 2015; Maggiotto and Pierson 1977; Mayer 2017; McGregor, Caruana, and Stephenson 2015; Medeiros and Noël 2014; Richardson 1991; Rose and Mishler 1998; Vlachová 2001).

Despite the small literature that conceptualizes negative partisanship as a social identity, it is important to recognize that many scholars are either skeptical of the existence of negative PSI or think that negative PSI is psychologically dissimilar to positive attachment. First, there are many theoretical reasons to expect negative and positive PID should function differently. For example, Caruana, McGregor, and Stephenson (2015: 774) point to the overwhelming evidence that human beings process negative and positive information (and emotions) differently:

> *"It has been established that a negativity bias exists, such that individuals react more strongly to negative than positive information; they are more likely to pay attention to it, more likely to remember it and likely to weight it more heavily when making decisions (Baumeister et al., 2001; Grabe et al., 2000; Newhagen and Reeves, 1992; Taylor, 1991). Many specific cognitive biases are related to this overriding power of negativity, including some that affect how people perceive economic consequences, such as loss aversion, endowment and anchoring biases (Baumeister et al., 2001; Carmon and Ariely, 2000; Kahneman and Tversky, 1979; Kahneman et al., 1990 Thaler, 1980)."*

Likewise, Zhong et al. (2008) argue that differences might stem from differences in the origin of negative and positive PSI: while the latter may develop to satisfy the basic human desire for inclusion (emphasizing similarity to other ingroup members) and distinctiveness (emphasizing distinctiveness from out-group members), the former is less likely to fulfill the need for inclusion.[30] Finally, a number of political scientists who think negative PID is worth studying are skeptical that it really is negative *identification*. For example, Rosema and Mayer (2020, p. 127) "doubt that for negative sentiments, self-categorization applies in the way that social identity theory presumes…" – a view that Oscarsson and Holmberg 2020: 26) seem to endorse. One possibility is that, despite some initial evidence to the contrary (e.g., Bankert 2020), negative identification may simply be expressions of out-party animus engendered by a strong positive in-group identity. Alternatively, even if negative partisanship is not just out-party animus (e.g., if it can happen in the absence of a

[29] The idea of negative partisanship may be particularly important in contexts in which citizens have overall low rates of positive partisanship, such as Latin America (Haime and Cantú 2022).
[30] See Appendix F.15.

strong in-party identity), that does not mean it is best thought of as a social identity (i.e., a positive identity with a negatively defined social group).

More generally, the theoretical ambiguity in which the whole concept of negative partisanship currently exists suggests that our conception of multi-partisanship would do well to allow for the possibility of negative PID and to try to measure negative and positive PID separately, which is one of the main goals of our measurement strategy for multi-party PSI. Further, to the extent that negative partisanship is real, it may be particularly helpful to study it in multi-party contexts.

First, if negative PSI is important separately from positive PSI, it is likely to be particularly important in situations in which voters are asked to evaluate the kinds of partisan coalitions between parties that are mostly absent in in two-party systems. For example, a voter who positively identifies with one coalition partner but negatively with another may find it hard to accept that coalition even when doing so would seem an otherwise rational way to achieve her policy or office goals.

Second, from a purely epistemological viewpoint, the overall study of negative PID may well benefit from the empirical leverage multi-party partisanship offers: While in two-party systems it is difficult to disentangle negative PSI from generic out-group affect, multi-party systems are much more likely to produce voters whose expressed negative attachments vary over the set of possible "out-parties" defined by their primary attachment(s). This provides more leverage to identify the distinct impact of negative attachments separately from the impact of positive ones on attitudes about out-parties. Of course, whether the theoretical concept of negative PSI will ultimately prove useful needs to be determined by appropriate data analyses, for which a focused measure of the concept for multi-parties is a necessary first step.

3. Measuring multi-party identity

In this chapter, we describe the survey items that we have developed to measure respondents' levels of positive and negative identification with each party in their system. We then explain how we use answers to these items in appropriate IRT models to produce estimates of θ_{ij} for each party. Along the way we discuss our methodological choices and demonstrate the robustness of our estimates to reasonable changes in these decisions. In the final section of this chapter, we address important questions about how to interpret our measures, developing new concepts for understanding concepts like a respondent's total partisan strength in the context of multi-party attachments. Due to space constraints, we relegate many methodological details and discussions to the online appendix (e.g., how long the survey took and the extent of respondent cooperation). To facilitate cross-referencing, we have numbered the appendix sections to correspond exactly to the section numbers and titles in this chapter (for this chapter only).

3.1 Developing the items in the survey battery

Over the last twenty years, there has been a great deal of progress in measuring the direction and intensity of partisan social identities (Bankert 2024; Huddy 2001, 2013; Huddy, Davies, and Sandor 2020; Schickler and Green 1997). These authors argue that unobservable psychological constructs like PSI are best measured with multi-item batteries constructed to cover the full range of the underlying trait. This approach is not new: for example, Miller and Shanks (1996) suggested the same, though it has taken a long time for the discipline to respond.

These works address the clear limitation of that traditional measures of partisan attachment "do not measure the full range of identify strength," despite the fact that "measuring gradations in social identity strength is crucial to identifying individuals who are most likely to vote for their party, practice defensive reasoning, and take political action" (Huddy, Davies, and Sandor 2020: 113). When properly constructed and tested, such multi-item scales provide more valid measure than single-item alternatives (e.g., Huddy and her colleagues 2001, 2013, 2015; Bankert 2024). We thus follow this literature in building a multi-item scale aimed at measuring an internalized sense of party identification, starting with items from previously tested batteries.

Further, since negative PSI has become an important research area in comparative political behavior, our battery aims to measure both positive and negative PSI. Usefully, item batteries designed to measure negative PSI for single parties have recently been introduced and so we can start with these (e.g., Bankert 2024). However, unlike previous efforts that have developed separate item batteries for negative and positive PSI, we construct a single set of balanced items that allows respondents to indicate either positive or negative (or no) attachment to that party. This design is essential for efficiently surveying attachments across many parties, replacing two separate question batteries for a large number of parties. As we show below, at the cost of a relatively innocuous assumption that one cannot be both positively and negatively attached to the same party, we can construct answer categories that capture the full range of positive and negative attachments.

Finally, while previous scholars have argued that multi-item batteries are necessary to measure the full range of partisan identify *strength*, we think a similar argument applies to its *direction*. All current multi-item measures of identity strength are applied only after filter questions that indicate the direction of partisan attachment (positive or negative). Not only does this tend to cut off the lower part of the attachment distribution (because respondents who do not pass the filter question are not asked about their partisan strength), but it also forecloses the possibility that individuals are attached to multiple parties, potentially painting an incomplete picture of a respondent's partisan identity.

To remedy this and to measure each respondent's full partisan identity, we administer our PSI battery for each party in the system. We then use the responses (in appropriate IRT models) to produce the partisan attachment vectors (with confidence intervals) illustrated in the introduction.

To summarize, our PIS measure aims to:

- Measure respondent *i*'s PSI for each of the J parties in their system (i.e., estimate $\vec{\boldsymbol{\theta}}_i$, which will have a separate estimate, θ_{ij}, for each party j).

- Estimate each θ_{ij} using a multi-item battery to ensure fine-grained coverage of the extent of PSI for party *j*.

- Include balanced items that capture both negative and positive identity while avoiding separate items for negative and positive identities (by relying on the assumption that one cannot be both positively and negatively attached to the same party).

- Design the survey and estimation methods so that the sign of θ_{ij} reflects positive or negative PSI for party *j*, while $\theta_{ij} = 0$ indicates no PSI for party *j*.

- Use a format that encourages respondents' attention, minimizes straight-lining or satisficing, and does not over-tax respondents.

- Follow modern best practices for designing such batteries (e.g., avoiding middle categories, randomizing answer categories and question order, and preventing item non-response while also providing appropriate guidance for those who might otherwise want to skip an item)

3.1.1 Constructing Items

Given the careful work other scholars have put into designing multi-item measures of PSI, there was no need to reinvent the wheel. Previous batteries have included items tapping important aspects of a social identity such as self-categorization, positive emotional experiences, relationship with other partisans, self-stereotyping and connection to other partisans.

Many of the items used in previous studies (e.g., Green, Palmquist, and Schickler 2002; Greene 2002, 2004; Ohr and Quandt 2012; Huddy et al 2015, Bankert et al. 2017) are based on Mael and Tetrick's (1992) Identification with a Psychological Group Scale (which applies to any group) and its evolution in subsequent studies (e.g., Ellemers et. al. 1999; Leach et al. 2008). For example, Bankert et al.'s (2017) 8-item partisan identity scale uses fairly direct translations of these items. Scholars have shown that scales based on these items have excellent scale properties and external validity. Thus, had we aimed only to measure positive PSI, we could have simply adopted one of these existing batteries. Some modifications were required, however, to produce a battery that allowed us to estimate both positive and negative identity from a single question battery. Even so, six out of the seven items we included in our surveys are quite close in Bankert et al. 8-item battery.

Our questions differ from previous efforts primarily in their response categories. To measure each respondent's level of positive or negative PSI with each party (θ_{ij}), where the sign of θ_{ij} indicates the direction (positive or negative, with a zero indicating no attachment to party j), we designed *balanced* items where the answer categories offer respondents two options that reflect increasingly positive identification and two that are increasingly negative, with no middle option.[31]

The five balanced items we ultimately used are shown in Table 1, as well as two other items in Table 2. We selected these seven items based on a pre-test (not reported here) in New Zealand in which we included items 1, 2, 4, 5, 6, and 7 (in Tables 1 and 2), plus two items tapping emotional reactions to the party's perceived performance – one from Huddy's 8-item scale asking how the respondent would feel if the party does badly in opinion polls, and another asking whether a respondent feels more proud or ashamed of the party's impact on society. The latter two items performed worse in our pre-test than the other items, revealing weakness in their coverage of extreme levels of (positive and negative) attachment. Thus, we added item 3, which adds strong extremal categories ("hate" and "love") to better capture very strong (positive or negative) attachments.

Table 1: The items included in our 5-item PSI battery

[1. Like]: "If the only thing you knew about a person was that they strongly identify with [party], how much would you like or dislike that person?" [reverse coded]	Like a lot	Like	Dislike	Dislike a lot
[2. Praise]: "When people praise [party] it makes me feel…" [reverse coded]	Very good	Good	Bad	Very bad
[3. Feeling]: "The word that best describes my feelings about [party] is…"	Hate	Dislike	Like	Love

[31] We follow Bankert, Huddy, and Rosema (2017) in excluding middle categories – based on the empirical finding that such categories tend to attract respondents who have weak but non-zero levels of the latent trait (e.g., Bishop 1990; Krosnick and Schuman 1988)

[4. Close]: "I identify most closely with people who…"	Strongly oppose [party]	Oppose [party]	Support [party]	Strongly support [party]
[5. Common]: "I have the most in common with people who…"	Strongly oppose [party]	Oppose [party]	Support [party]	Strongly support [party]

The short labels provided after the question numbers refer to these items below.

Besides the five items in Table 1, we initially included two other items that have been mainstays of the previous literature and directly tap the self-categorization aspects of social identity (as we would argue item 4 also does). Their answer categories, however, are not balanced in the same way as those in Table 1. For example, strong disagreement with the "my party" statement in item 6 may indicate non-partisanship (with respect to this party) rather than negative identification. Likewise, not feeling connected to a person signals an absence of positive emotion, not the presence of a negative one.

Table 2: The items included in the survey, but excluded from our 5-item PSI battery

[6. My Party]: "When I think of [party], I think of it as 'my political party'." [reverse coded]	Strongly agree	Agree	Disagree	Strongly disagree
[7. Connected]: "When I meet a supporter of [party], I feel…" [reverse coded]	Strongly connected to that person	Somewhat connected to that person	Not very connected to that person	Not at all connected to that person

Given the lack of balance in these items (and no clear way to make them balanced), we ultimately did not include them in the final measure we use in the rest of the manuscript. The potential cost of this is omitting an indicator of self-categorization and belonging that has often been included in similar scales (the "my party" item). However, we show in Appendix A that this is not a problem. There we provide a variety of analyses that show 1) the "my party" item is very strongly correlated with the five items we ultimately use; 2) an IRT analysis shows that this item (and item 7) is not essential for a reliable, unidimensional partisanship scale; and 3) individual-level estimates of attachment to each party are largely unchanged when these items are added to the 5-item index and scaled together.[32]

[32] However, because these measures are not balanced, adding them to the scale complicates the interpretation of what an estimated attachment of zero means.

The reason we think our five-item index can safely exclude the "my party" item is that it includes other items that tap self-categorization. Specifically, item 4 explicitly invokes the concept of identity with the group. In support of this interpretation, of all the items in our batteries, item 4 has the highest correlation with the "my party" item (.64).

Given the close correspondence between answers to the "my party" item and our index based on the five balanced items, we think the latter is a useful measure of psychological attachment, including a sense of belonging. That said, in Appendix G, we discuss a number of alternative batteries that may be of use to researchers who need fewer items or do not need a balanced (negative-positive) measure of multi-party partisanship.

Finally, the "my party" item may be particularly useful for persuading otherwise skeptical readers that people really can have psychological attachments to multiple parties who periodically compete with one another in elections. The "my party" item directly asks whether respondents think of a given party as *their* party, directly tapping into feelings of belonging. If multi-party identification were implausible, our respondents should balk at telling us, in such a direct way, that more than one party is "their" party.[33] Yet, Figure 2 shows they do not: 77%, 73%, 73%, and 58% of our respondents in Denmark, the Netherlands, Sweden, and Italy, respectively, indicate that more than one party is "their party." Moreover, these declarations are not at all capricious; they are structured exactly as we would expect. For example, respondents are much more likely to say that two parties are "my party" if they are on the same side of the left-right divide.

[Figure 2 here]

3.1.2 Measurement of negative PSI

Tables 1 and 2 explain how we modified common positive PSI items to produce measures of both positive and negative PSI. We are not the first, however, to build a negative PSI scale by modifying existing positive items (e.g, Bankert, Huddy, and Rosema 2017; Bankert 2020, 2024). The key difference is that we employ a single symmetric battery administered to all respondents for each party, while previous measures used two separate batteries for positive and negative PSI, assigned via filter questions.

In contrast, we do not need filter questions. Filter questions – such as vote choice – while common in prior work, are problematic (as the scholars who use them are very much aware). They limit the range of identity strengths that can be identified, since they necessarily make assumptions about who gets the identity battery and who does not. A vote choice filter, for example, will certainly miss some lower-level identifiers who did not vote or voted for a different party or even higher-level

[33] Obvious examples in which people identify with more than one social group that (at least sometimes) compete include identification with different countries (e.g., immigrants), families (e.g., one's own family and one's partner's family), and sports teams (e.g., one's hometown team and the team where one currently lives).

identifiers who have temporarily deviated from voting for the party (or who did not vote). This may be particularly problematic in "low partisanship" countries (Guntermann 2019). Simply asking the battery for all parties for all voters removes the need for a directional filter question and minimizes an increasingly relevant kind of social desirability effect -- the hesitation to identify oneself as a partisan -- since respondents are never directly asked to do so (Baker and Renno 2019; Klar and Krupnikov 2016).[34]

The need for a directional filter question before assessing partisan strength is even more challenging for negative PSI. The few scholars who have done so have usually turned to a question (commonly included in European surveys) that asks respondents to identify parties for which they would never vote. This filter introduces the same problems that Huddy, Davies, and Sandor (2020) bemoan for positive PSI. Our battery avoids the need for this kind of filter entirely and captures negative and positive identity in one battery.

What is the cost we pay for that? Besides survey length, our balanced battery assumes that one cannot hold a negative and positive PSI to the same party. Technically, this is not a requirement of Bankert's positive and negative PSI batteries. Both batteries *could* be given to each respondent, but practically this is not usually done. Instead, filter questions are used to determine which respondent gets which battery (but with some respondents getting both in Bankert's 2024 book).[35]

3.1.3 The surveys

We fielded a pilot survey in New Zealand in early November of 2021. After adjusting the instrument based on pilot results, we fielded surveys in Sweden, Denmark, Italy, and the Netherlands from mid-December 2021 to early-January 2022. In November 2024, we conducted an additional survey in the UK that included our PSI battery and additional party performance items. Samples of approximately 1100-1200 respondents per survey were recruited from Qualtrics Panels (Lucid/Cint for the UK) and drawn to match census targets on age and gender. Respondents were part of large panels and received (randomly) an invitation that did not mention politics. Respondents were compensated by the survey vendor only if they passed tests quality checks for speeders and straight-lining for matrix-style questions (see Appendix A3.1.3 for details about the survey format and respondent experience).

Table 3: Comparing levels of (single party) identification and political interest in our surveys and the latest CSES surveys

	PID	Interest

[34] Huddy, Davies, and Sandor (2020: 111) point out that some sub-populations seem to be particularly hesitant to declare a partisan direction (e.g., American Hispanics) despite acting very much like partisans in other ways.

[35] See appendix F.16.

Country and date of CSES survey	CSES	Our surveys	Difference	CSES	Our surveys	Difference
Denmark 2019	79.6	89.7 (~2022)	**10.1**	75.9	84.3	**8.4**
Italy 2018	68.1	80.5 (~2022)	**12.4**	59.3	74.3	**14.9**
Netherlands 2021	74.4	85.2 (~2022)	**10.9**	68.7	80.9	**12.2**
Sweden 2018	83.5	90.4 (~2022)	**6.9**	67	76.7	**9.7**
UK 2019	76.8	86.2 (~2025)	**9.4**	73	81.4	**8.4**

The PID column indicates the proportion of respondents who answered the question and said they were close to one party either after the first prompt or after the follow-up, "nudge", question (our PID question was identical to CSES). The interest column is the proportion of respondents (who answered the questions) who reported very interested or somewhat interested in politics (our political interest question was identical to the question in CSES). Our original surveys were conducted in November-December 2024 for the UK and November-December 2021 for the other four countries. These proportions are all from large samples, so standard errors are tiny and all the reported differences are statistically different from zero.

Table 3 compares average levels of partisanship and political interest in our samples to the most proximate CSES surveys for each county.[36] Reported rates of (traditional) party identification in the CSES surveys differ from our original surveys by about ten percentage points. Differences in the timing of the surveys might explain some of these differences, but the differences are all positive and not associated with the time between the CSES surveys and ours. Further, our samples show higher rates of party identification and political interest than the CSES probability samples.[37] This overrepresentation of politically engaged respondents is not surprising, given that we are using opt-in online panels in which some politically uninterested respondents may drop out once they discover the survey is about politics.

That said, these differences are modest. Indeed, since our primary objectives in this Element can be accomplished by focusing on ***within*** country comparisons (and usually within survey), we do not attempt to adjust samples to correct for this issue.[38] Instead, we simply point out this limitation here and avoid making strong conclusions about quantities that might be influenced by this (relatively mild) oversampling of politically interested respondents.

[36] See Appendix F.17.
[37] See Appendix F.18.
[38] See Appendix F.19.

3.2 Estimating and evaluating the partisan attachment vector

In this section, we briefly summarize the methods used to estimate each respondent's partisan identity – i.e., the party attachment vectors illustrated in the introduction.

Treating party attachment as a latent trait (see Bankert, Huddy, and Rosema 2017; Green, Palmquist, and Schickler 2002), we use an Item Response Theory (IRT) approach to scale building. IRT models the relationship between a set of item responses and the latent "ability" of interest, without assuming that each item is equally "difficult" or that it contributes equally to the latent trait. It also assesses the performance of individual items at discriminating various levels of the latent trait. Since we build our attachment scale using ordered and polytomous responses, we use Samejima's (1969) graded response IRT model and estimate each respondent's attachment to each party (θ_{ij}) using expected a posteriori (EAP) methods (summarized in an estimated attachment vector $\theta_{ij} \in \vec{\boldsymbol{\theta}}_i$).

To save space in the main text, we relegate most technical details the online appendix (Section A3.2), including a discussion of whether to estimate the IRT models pooled over parties or for each party separately (Appendix A3.2.1), which respondents to include (Appendix A3.2.2), and evidence of internal validity and overall scale performance (Appendix A3.2.3). In summary, we estimated IRT models for each party separately rather than pooling within country (but little depends on this choice). We excluded respondents who reported themselves unfamiliar with the party in question from these party-specific IRT models (but included them in the subsequent substantive analyses, assigning $\theta_{ij} = 0$ for the unfamiliar party). Finally, all scales pass the standard IRT and Classical Test Theory checks for scale coherence, reliability, and uni-dimensionality.

[Figure 3 here]

Figure 3 illustrates how this scale differs from one constructed as a simple five-item index. Each panel shows the range of estimated values of θ_{ij} that are associated with each possible values of a simple additive index (or "Sum Score") constructed from the same items (given 5 items, with answer categories coded 0-3, the minimum sum score is 0 and the maximum is 15). As shown, IRT methods produce a range of θ_{ij} values for each sum score, while producing values that have a clear and strong positive relationship with it.

More generally, IRT models the probability of a specific response as a non-linear (in our case, logistic) function of both the respondent's underlying unobservable attachment θ_{ij}, and a set of item parameters (the item's difficulty and discrimination). The first stage of the process is to estimate the item parameters using all data, integrating the unknown θ_{ij}'s out of the likelihood function. The second stage takes those item parameters as given and uses each respondent's responses across all items to find the Maximum Likelihood estimate of θ_{ij} that would be most likely to produce that specific response pattern. This fundamentally differs from a simple additive scale, which assumes all items are equivalent and merely sums the item scores. In an additive scale, a high valued response to a highly discriminating item contributes the same as a less discriminating item. In IRT, a high valued response to a highly discriminating item contributes more weight to the ultimate score.

3.3 Interpreting $\vec{\boldsymbol{\theta}}_i$

How should we interpret each respondent's partisan identity vector, $\vec{\boldsymbol{\theta}}_i$? In this section, we first defend the idea that positive values of θ_{ij} indicate positive attachment to party *j*, negative values indicate negative attachment, and zero indicates unattached. Likewise, larger (absolute) values of these quantities indicate stronger attachment – directly analogous to single-party measures of partisanship and negative partisanship. However, we also address two less straightforward questions that comes up only in the case of multi-party identity: (1) How should we understand total identity strength when these identities may be spread across multiple parties (and be a mix of positive and negative identities); and (2) How should we conceptualize the strength of a voter's level of *co-partisanship* between any two parties?

3.3.1 How can we interpret the sign of θ_{ij}?

By construction, our estimated θ_{ij}'s are distributed standard normal, with $\theta_{ij} = 0$ representing the level of the underlying trait (negative to positive PSI) for the average respondent for party j. The first question we ask is: under what conditions we are justified in interpreting $\theta_{ij} = 0$ as being not attached (whether positive or negative) to party j? In addition, can we interpret the sign of θ_{ij} as indicative of positive vs. negative PSI? Usefully, we can use the estimated item parameters from our IRT model to examine the validity of these interpretations.

The answer to these questions depends on how we constructed the items (and answer categories) that contributed to our scale and whether the items performed as expected. With respect to item design, each item has four diametrically opposed answer categories, so a non-partisan should avoid both the extreme answer categories for each item. We deliberately omitted a neutral category, which would otherwise attract nonpartisans. This reflects best practices, as middle categories tend to attract responses from individuals who are not actually neutral on the question, but only weakly so.[39] We also require respondents to answer every item. Given this, it is important to think about how a true non-partisan (with respect to party j) would answer the battery to indicate that. Our assumption, following the empirical conclusions of O'Muircheartaigh, Krosnick, and Helic (2001), is that such respondents will randomize between the two moderate categories that define the qualitative division between negative and positive PSI.

Usefully, the item parameters from the graded IRT model provides information we can use to check the veracity of this assumption, and whether $\theta_{ij} = 0$ indicates not just average (positive to negative) partisanship, but a lack of positive or negative partisanship with respect to party j.

Figure 4 plots the Category Response Curves (CRC) for an item highly consistent with interpreting $\theta_{ij} = 0$ as non-partisanship and its sign as indicative of positive or negative PSI. The x-axis represents θ_{ij}, and the y-axis indicates the probability of choosing each answer category. Since respondents had to choose one option, the probabilities for all categories at any point on the x-axis sum to one.

[39] See Appendix F.20.

[Figure 4 here]

Reading up from 0 on the x-axis, a respondent with $\theta_{ij} = 0$ has a 50/50 chance of picking one of the moderate categories for this item – exactly the assumption we make about how such individuals will answer. Likewise, the probability of choosing the most extreme negative category is 0 for everyone with a positive θ_{ij} and the probability of the moderate negative category approaches 0 quickly as θ_{ij} gets more positive. The situation is symmetric for negative θ_{ij}'s. If all CRC plots for a given party looked like Figure 4, we should be quite comfortable interpreting $\theta_{ij} = 0$ as non-partisanship with respect to party *j*.

Figure 5 provides all five CRC plots for another example, the Swedish Moderate Party, which is quite typical of our party-specific results. Each item is quite close to the ideal illustrated in Figure 4, though one item ("like") deviates somewhat.

A more general approach to this issue uses the item parameters underlying the CRC plots to calculate, for each item for each party, the probability of choosing answer category 3 or 4 when a respondent's $\theta_{ij} = 0$. If this probability is 50%, the respondent is equally likely to choose a negative (1 or 2) or positive (3 or 4) category. Thus, if this probability is near 0.5 for all our items and parties, we have a strong justification for interpreting the sign of θ_{ij} (and a value of 0) as meaningful in the ways described above.[40]

[Figure 5 here]

This is, indeed, generally what we find for our five items. Figure 6 provides the estimates of this probability for all items across all parties (using party-specific IRT models). Given that 85% of these probabilities are within 10 percentage points of 50%, we think it is safe to interpret the sign of θ_{ij} as indicative of positive or negative PSI, with values near 0 indicating non-partisanship with respect to party j.

[Figure 6 here]

A final issue of interpretation concerns how we should think about individuals who evidence a negative attachment to a party. As previewed, the literature on negative partisanship is hardly settled, with many scholars doubting that positive and negative partisanship are two sides of the same coin. Instead, some argue that what the literature has identified as negative partisanship is simply out-party animus stemming from positive identification with some other party. In this view, the negative "attachments" that our measure produces would reflect the usual biases against out-groups rather than identification with a negatively-defined social group. If this is the case, negative identity should not "work" the same as positive identities. For example, stronger negative attachments should not be associated with greater levels of partisan rationalization in favor of a negatively-defined out-group. In Section 5.2 we explore various partisan rationalization and find

[40] We also require that our extremal categories are well behaved in that there is very low probability that a respondent with $\theta_{ij} = 0$ would choose them, which holds for all our items.

mixed results. This suggests that partisan rationalization does not always work the same way for positive and negative partisans, giving some credence to the idea that they may be psychologically distinct.

Other evidence consistent with this view comes from examining whether a positive partisan identity is a necessary condition for negative identification (as measured by our scale). If negative identification is just out-party animus, we should not see negative identification in the absence of positive identification.

Figure 7 provides data from Sweden that are relevant to this question (other countries produce similar results). The points in the graph show, for each respondent, the value of their *largest* and *smallest* estimated θ_{ij}. In the left panel, the five respondents in the lower-left quadrant are the only ones in which all their estimated θ_{ij} (across all parties) are negative. In contrast, all cases in the upper-left quadrant have at least one negative and one positive party attachment. As we move toward the northwest corner of that quadrant, we find individuals who are both strongly positively attached to one party and strongly negatively attached to another – consistent with the idea that positive attachment to an in-party begets negative attachment to an out-party. The clear linear trend in this quadrant also supports this view.

[Figure 7 here]

That said, these data are not definitive. It could be that while negative identifications are more likely among those who also positively identify, the latter does not *require* the former. At this point, it seems prudent to be cautious about what exactly our negative partisanship measure captures and to interrogate whether it has the same effects on downstream attitudes and behaviors expected from positive partisan attachments (e.g., evidence of partisan rationalization in favor of the negatively defined social group).

Finally, in Appendix G we evaluate several alternative ways to estimate multi-party attachments, including, for researchers uninterested in negative partisanship, a version of our battery that does not require symmetric items and can include unbalanced items (like the "my party" item) that we excluded.

3.3.2 How to conceptualize and measure the total strength of partisanship for multi-party partisans?

Given that $\vec{\boldsymbol{\theta}}_i$ is a vector in J-space and the origin of that space at 0 is interpretable as a pure independent, the most obvious measure of total partisan strength across multiple parties is the length of $\vec{\boldsymbol{\theta}}_i$, usually denoted $||\vec{\boldsymbol{\theta}}_i||$.[41] We can also focus on only the positive (or negative)

[41] This length is the sum of the squares of the component θ_{ij}'s.

contributions to the vector by recoding negative (or positive) θ_{ij}'s to 0 before calculating the length of the vector (these zeros add nothing to this length).

More formally, define *i*'s partisan identity as $\vec{\boldsymbol{\theta}}_i$, the strength or her identity with respect to party j as θ_{ij}, and the overall strength of her identity as the length, or magnitude, of the vector $||\vec{\boldsymbol{\theta}}_i||$. Likewise, define respondent i's positive partisan identity, which we will label $\vec{\boldsymbol{\theta}}_i^+$, as equal to $\vec{\boldsymbol{\theta}}_i$ with all its negative elements set to zero, and the strength of respondent i's positive partisan identity as $\left\|\vec{\boldsymbol{\theta}}_i^+\right\|$. Analogously, define negative partisan identity, $\vec{\boldsymbol{\theta}}_i^-$, as $\vec{\boldsymbol{\theta}}_i$ with all its positive elements set to zero, and the strength of i's negative partisan identity is $\left\|\vec{\boldsymbol{\theta}}_i^-\right\|$.

Figure 8 illustrates these concepts for five voters (voters A, B, C, D, and E) and two parties (Social Democrats and Green) in a two-dimensional "partisan identity space." Voters whose partisan identities fall on one of the two axes (e.g., voter B) are "pure" single-party partisans, while those whose partisan identities do not fall on these axes are multi-party partisans. Voter A's partisan identity vector starts at the origin {0,0} and extends toward the upper right to the point labelled $\vec{\boldsymbol{\theta}}_A$. It captures the direction of voter A's full partisan identity (i.e., equal attachment to both parties), the strength of her attachments to each party (the party specific components of $\vec{\boldsymbol{\theta}}_A$), and her overall identity strength (i.e., $||\vec{\boldsymbol{\theta}}_A||$).

[Figure 8 here]

Comparing voter A to voter B illustrates an important aspect of defining total partisan identity in this way. Notice that voters A and B are equally attached to the Social Democrats (level *w*) and that voter A is additionally attached (by an equal amount t = *w*) to the Greens. Voter A's identity strength is $\sqrt{w^2 + t^2} = \sqrt{w^2 + w^2}$, while voter B's is only $\sqrt{w^2}$. In the usual single-party setting, both A and B would be considered equally partisan, while in this multi-party setting, voter A is "more" partisan. Using the vector length as a measure of total partisan strength, we are implicitly adopting a specific function for how total partisan strength grows with additional partisan attachment; and this function is conservative relative to plausible alternatives (like summing the absolute values of each θ_{ij}). One can see this by noticing that if *w* were equal to 1, so that voter B's partisan identity strength would be $\sqrt{1^2} = 1$, voter A's partisan identity strength would only be $\sqrt{1^2 + 1^2} \approx 1.41$, despite being equally attached to both parties. Had we instead simply summed up the (absolute) values of A's attachments to each party, we would have instead found A's partisan identity strength to be equal to 2. Thus, the vector length displays a kind of diminishing marginal return to increasing attachment that the summation measure does not – and, while it is ultimately an empirical

question which measure will prove more useful in capturing the concept of total partisan identity, from a theoretical perspective this feature seems quite plausible.[42]

3.3.3 How to conceptualize and measure the strength of co-partisanship

An important new concept that emerges when one allows for multi-party partisanship is the strength of co-partisanship for any two parties. Building measures of this for each respondent and party-dyad allows us to test hypotheses about the causes and consequences of co-partisan identity – for example, does it strengthen when the parties in question join the same coalition?

Figure 9 illustrates our approach to conceptualizing (and measuring) co-partisanship between two parties *j* and *k*.[43] It depicts three voters with different strengths of co-partisanship. Since we focus here on each two-party sub-space of the whole J-party partisan identity space defined above, we adopt the notation $\theta_{i(jk)}$ to mean the part of voter *i*'s partisan identity vector that applies to parties *j* and *k*. The illustration is drawn so that $||\theta_{A(jk)}|| = ||\theta_{B(jk)}||$, i.e., the total partisan strength for these two parties is the same for voters A and B.

[Figure 9 here]

The fact that none of the vectors $\vec{\boldsymbol{\theta}}_{i(jk)}$ fall on the x or y axes means that all three voters (A, B, and B) have some attachment to both parties. But how do their levels of co-partisanship compare to one another? Intuitively, a voter who has more (less) attachment (than another voter) to either of the two parties should be considered more (less) co-partisan. Another way to say this is that as the total partisan strength for the two parties increases, the strength of co-partisanship should also increase. We can see this in Figure 9 by comparing voters A and C. While they identify equally with parties *j* and *k* (i.e., both vectors are at angle of 45 degrees), voter C is less attached to both parties than voter A ($||\vec{\boldsymbol{\theta}}_{A(jk)}|| > ||\vec{\boldsymbol{\theta}}_{C(jk)}||$). Thus, voter C should be thought of as being less strongly co-partisan (for parties *j* and *k*) than voter A.

What about the comparison between voters A and B? While voter A and B's partisan identity strengths for this party-dyad are equal ($||\vec{\boldsymbol{\theta}}_{A(jk)}|| = ||\vec{\boldsymbol{\theta}}_{B(jk)}||$), we argue that we should ***not*** conclude they have the same level of co-partisan attachment to these two parties: voter B identifies more closely with party *j* than with *k* – a significant asymmetry compared to voter A (who has no asymmetry). In our view, the theoretical concept of strength of co-partisan attachment is more useful if it accounts for this kind of asymmetry, because there are many interesting questions in which we would not expect voters A and B to be equivalent. For example, when evaluating a potential governing coalition between the parties, we would likely expect differences between

---

[42] See Appendix F.21.

[43] The concept and measures of co-partisanship that we develop in this section (and that we use in Chapter 4) apply to positive co-partisanship and so are conceptualized and calculated using $\vec{\boldsymbol{\theta}}^{+}_{i(jk)}$. For notational simplicity, where it will not lead to confusion, we drop the "+" notation in this rest of this section.

voters A and B due to the asymmetry in their identification with each party: voter A might accept more policy compromise between the parties for the coalition to form, while voter B would likely be happier with less compromise from party *j* relative to *k*.[44]

We can use these concepts to develop measures of co-partisanship that depend on both the overall strength of partisanship for any two parties and the level of symmetry in these identifications. In Chapter 4, we use a measure of co-partisanship built as follows. First, we use our estimated $\theta_{i(jk)}$ to calculate the angle $\lambda_{i(jk)}$ associated with the voter's attachment vector for a given party pair. Next, we transform this angle (deterministically) into a variable ranging from 0 to 1, where 1 indicates a 45 degrees angle and 0 indicates an angle of either 0 or 90 degrees, with intermediate angles varying accordingly (according to a linear function) – see details in Appendix C.[45] Intuitively, this captures the extent to which the partisan sub-vector for any two parties is at 45 degrees (indicating equal identity to both parties). We then multiply this by $||\vec{\boldsymbol{\theta}}_{i(jk)}||$, our measure of total partisan identity strength toward the two parties.[46] The resulting quantity increases when the voter's identity with one or both of the parties strengthens but decreases as these identities become less balanced between the two parties.

[44] This means that for any two parties, voter i's total partisan attachment to those two parties can differ from the strength of their co-partisan attachment to the same two parties – emphasizing that these are conceptually distinct concepts.

[45] Other functions could be used. Lacking any theoretical guidance on this functional form, starting with a linear function seems reasonable.

[46] See Appendix F.22.

4. Describing patterns of multi-party identity in in Denmark, Italy, Sweden, the Netherlands, and the UK

With estimates of each respondent's full multi-party identity in hand, we now explore the extent and nature of multi-party identity in our sample of countries and compare them to traditional, maximalist PID measures. First, we briefly describe broad features of the multi-party identity revealed by our estimates: how common are multi-party identities relative to single-party ones? And how do our estimates of the proportion of independents within each survey compare to traditional measures?

Next, we summarize a set of analyses that explore the structure of multi-party identities. For example, what does knowing the extent of (positive or negative) attachment to one party tell us about (positive or negative) attachment to another? One way we do this is to examine the extent of co-partisanship (as defined in Chapter 3) for different party pairs.

Finally, given space constraints, many analyses in this chapter focus on positive partisanship, primarily to fairly compare multi-party partisanship to traditional measures (which do not include negative partisanship). However, we bring in negative partisanship more fully in Chapter 5 when we examine how multi-party partisanship impacts downstream attitudes and behavior. However, given the theoretically ambiguous status of negative partisanship as a concept (as discussed in sections 2.5 and 3.3.1), these analyses should be taken as preliminary. Negative partisanship deserves more complete theoretical and empirical exploration than space here allows.

4.1 How does $\vec{\theta}_i$ compare with traditional measures of PID strength and direction?

In this section, we compare our estimates of each respondent's complete partisan identity, $\vec{\theta}_i$, to traditional measures of partisan direction and strength. Specifically, we ask:

- When the traditional PID measure identifies a party that a respondent feels closest to, do we identify the same party as the one with the highest estimated θ_{ij}? If not, is our estimate of θ_{ij} for the party close to the largest?[47]

- When the traditional PID measure indicates a respondent is non-partisan with respect to a particular party (whether or not they are a traditional partisan of another party), how often do we identify significant partisan attachment to that party?

- When the traditional measure identifies a respondent as an independent, how often do we do the same and what issues arise in identifying independents using our measure?

[47] This second question is addressed in the online appendix.

Our data are organized at the respondent-party level, allowing direct comparison of our θ_{ij} estimates with traditional measures of PID (where PID is coded 1-3 for increasing strengths and can have positive value for only one party per respondent).

We begin by examining, in Figure 10, how our estimates of θ_{ij} correlate with traditional PID in each country. There is a very strong correspondence between our θ_{ij} and the traditional measure of PID strength, and relatively few cases have traditional PID > 0 but a negative θ_{ij}.

[Figure 10 here]

This correspondence, however, does not hold when the traditional PID measure is zero. Examining this more closely, Figure 11 shows the distribution of our θ_{ij} estimates for respondent-party dyads for which traditional PID for that party is 0. While this distribution is concentrated on zero, the variability of θ_{ij} for these traditional non-partisans (with respect to party j) is substantial.

[Figure 11 here]

This figure reveals that for many respondent-party pairs, the traditional assignation of no partisanship (with respect to party *j*) hides significant negative PSI. That is not surprising since traditional PID does not measure negative attachment. It is also the case, however, that many respondents answer our attachment batteries in ways that suggest significant *positive* attachment to parties for which their traditional PID is zero. We saw one example of this in the introduction to this Element: a Swedish respondent reported a traditionally measured attachment to the Left Party, which matched our high estimate of $\theta_{i,\text{LeftParty}}$, but was equally attached to the Greens, an attachment the traditional measure assumes is 0.

Another useful way to look at these traditional non-partisans (with respect to party *j*) is to examine the percentage of these traditional non-partisans whose θ_{ij} is statistically different from zero. We do this in Table 4, revealing that for all countries except the UK, about half of these non-partisan respondent-party pairs (for whom traditional PID measures indicate no identity with the party) have statistically significant attachments to the party – positive or negative – using our measure.

Table 4: Proportion of θ_{ij}'s that are statistically different from zero among traditional non-partisans (with respect to party j)

	Proportion of significant θ_{ij}'s, traditional non-partisans	Proportion of significant θ_{ij}^{+}'s, traditional non-partisans	Proportion of significant θ_{ij}^{-}'s, traditional non-partisans
Denmark	45.6%	45.1%	44.4%
Italy	53.8%	47.3%	51.1%
Netherlands	46.1%	47.9%	47.0%
Sweden	48.9%	56.2%	55.3%
UK	38.4%	27.5%	39.8%

As we discuss below, the UK's lower level of partisanship using our measure our measure (but not traditional measures) likely reflects depressed levels of co-partisanship in the UK relative to other countries – a result consistent with our hypothesis from Chapter 2 that the historical rarity of coalition cabinets would lead to less co-partisanship (In Section 4.3, we test this hypothesis with a direct measure of the extent of co-partisanship).

One finding we did not anticipate, but that may also be attributable to the UK's (near) lack of coalition government, is its relatively higher ratio of negative partisan attachment to positive attachment compared to other countries. If the impact of coalition government on co-partisanship comes from the necessity of inter-party cooperation, it stands to reason that this impact would be felt more in the development of positive than negative co-partisanship.

We can also disaggregate the results in Table 4 by examining the correspondence between the traditional PID strength and our estimates for each party. Figure 12 focuses on Sweden (others in Appendix C), showing how often we estimate a significant (positive) attachment to each party when traditional measures indicate the respondent identifies, or not, with the party.

[Figure 12 here]

This table (and the corresponding ones in Appendix C) make it clear that, for the vast majority of respondent-parties, both measures tell the same story.[48] That is, in about 90% of the cases where the traditional measure identifies a partisan, we estimate a statistically significant θ_{ij}.[49] Likewise, in

[48] See Appendix F.23.
[49] See Appendix F.24.

nearly 70% of the cases where the traditional measure indicates that a respondent is not a positive partisan of a particular party, our estimate agrees (θ_{ij}^{+} is not statistically different from zero). However, we often (about 30% of cases for each party) find respondents who have a statistically significant attachment to a party that is not revealed by the traditional measure. As we discuss below, almost all these cases come from multi-party partisans who are traditionally non-partisan with respect to party *j* because they named another party in the traditional maximalist question; but expressed multiple identities when given opportunity to do so in our battery. This ultimately contributes respondent-party observations to both the upper and lower right cells in Figure 12. For each party except the Centre Party, our measure fails to identify traditional partisans (with respect to party *j*) only about 10% of the time.

One drawback of this analysis is that it only checks if we estimate a significant partisan attachment for the *same* party that the traditional measure does. It could be, however, that voters with multiple attachments vacillate overtime with respect to which identity is strongest (and so also vacillate in how they answer a traditional PID question). Thus, we explore, for cases in which the traditional PID question and our estimates disagree on the party to which a respondent is most attached, whether our estimated level of attachment to the traditional PID party is "close" to the strongest one. While we relegate the details to Appendix C, the short answer is a clear "yes" – when a respondent identifies with a party in the traditional question, our estimate of her attachment to that party is almost always either the largest or almost the largest (in a precises sense defined in Appendix C).

4.2 How many independents, single-party, and multi-party partisans?

In the last section, we examined how our estimates corresponded with traditional PID measures (that only allow identity with one party). Here, we broaden our view to concepts that require consideration of the whole vector of identities θ_i in order to examine what it tells us about the proportion of different voter types that can only be defined by the whole vector: pure independents, single-party partisans, and multi-party partisans.

We identify a partisan with respect to a given party if their θ_{ij} for that party is statistically different from zero. We define an independent as any respondent for which none of her θ_{ij}'s are statistically different from zero. Similarly, a single-party partisan is a respondent who has only one statistically significant θ_{ij}. Multi-party partisans can be defined analogously.

One problem with applying these criteria naively is that each requires, by definition, multiple hypotheses tests – one for each of the J parties in a country. Thus, we apply a Holm correction in our analyses to account for multiple hypotheses testing. [50]

[50] The Holm adjustment for multiple comparisons aims to control the family-wise error rate (FWER) by sequentially comparing each p-value to a progressively adjusted significance level based on its rank among all comparisons.

Finally, to make comparisons with traditional single-party measures (which do not contemplate negative partisanship), we again focus on only positive partisanship by re-coding negative estimates of θ_{ij} to 0 and labelling them θ_{ij}^{+}).

Table 5 provides our estimates for each of the five countries. The first thing to notice is that about 50% of our respondents identify with more than one party. [51] The number is somewhat higher in Sweden (60.7%) and much lower in the UK (only 27.5%), consistent with our hypothesis that historical coalition governance fosters co-partisanship. Looking at the number of parties with which our multi-party partisans positively identify, we see that this is usually limited to a few parties: in Denmark and the Netherlands, the rate largely drops off after two parties, while in Italy and Sweden the largest drop off is after three parties; and in the UK this happens after one party.

Table 5: Levels of independence, single-party partisanship, and positive multi-party partisanship [52]

Number of θ^{+} that are statistically different from zero

	0	1	2	3	4	5	6	7	8	9	Total
Denmark	23.3% 255	27.4% 300	25.5% 279	11.7% 128	8.4% 92	2.7% 30	0.6% 7	0.0% 0	0.0% 0	0.4% 4	100% 1095
Italy	28.6% 326	20.4% 233	21.8% 248	20.1% 229	5.7% 65	2.1% 24	1.3% 15				100% 1140
Netherlands	21.2% 213	33.0% 332	20.0% 201	14.4% 145	7.0% 70	3.1% 31	0.8% 8	0.6% 6	0.0% 0	0.1% 1	100% 1007
Sweden	15.8% 159	23.5% 237	25.8% 260	19.9% 200	10.3% 104	3.4% 34	1.1% 11	0.0% 0	0.2% 2		100% 1007
UK	35.4% 417	37.1% 437	16.9% 199	6.2% 73	1.7% 20	1.0% 12	0.6% 7	1.1% 13			100% 1178

Significance values are determined using the Holm correction for multiple comparisons. For comparability with traditional partisanship, this table uses θ^{+}, which sets all negative values of θ to zero.

[51] See Appendix F.25

[52] To avoid confusion, it is important to keep in mind that the numbers for the percentage of independence (and partisans, of any strength) in this table are calculated using our PSI battery for each party, while the numbers reported earlier, in for example, section 3.1.3, are our estimates of the percentage of partisans using the traditional measure.

Next, the number of pure independents we find using our question batteries (applying the Holm method for determining statistical significance) is quite close to that found in the CSES for each country (see the column labelled "0"). Using the data reported in Table 3 in Section 3.1.3, the rate of pure independents in the CSES (from years closest to fieldwork of our surveys) is 20.4%, 31.9%, 25.6%, 16.5%, and 23.2% for Denmark, Italy, the Netherlands, Sweden, and the UK respectively.[53]

Finally, our estimates of the proportion of single-party partisans are between 20-30% in each country, except for the UK. The higher rate in the UK (again) fits our expectation that a history of frequent coalition governance is an important way in which partisans turn into co-partisans. The UK's low rate of co-partisanship is especially apparent when summing the first two columns of Table 5 and subtracting it from 100, indicating the proportion of co-partisans in each country: 49.3% in Denmark, 51.0% in Italy, 45.8% in the Netherlands, 60.7% in Sweden; but only 28.5% in the UK.

Now that we have established how to identify various voter types using our estimates of the partisan attachment vector, we can ask how closely our categorizations (independents, single-party partisans, or multi-party partisans) match these same designations using the traditional measures. In Table 6, we disaggregate the numbers from Table 5 by responses to the traditional maximalist measure for Denmark, illustrating the pattern found in all our other countries (see Appendix C).

Table 6: Correspondence between traditionally measured independence and single-party partisanship and our estimates of $\boldsymbol{\theta}^{+}$in Denmark

	Number of elements of $\boldsymbol{\theta}^{+}$that are statistically different from zero										
	0	1	2	3	4	5	6	7	8	9	Total
Trad. Partisan	20.1% 197	27.2% 267	27.5% 270	12.4% 122	9.0% 88	3.0% 29	0.5% 5	0.0% 0	0.0% 0	0.4% 4	100% 982
Trad. Independent	51.3% 58	29.2% 33	8.0% 9	5.3% 6	3.5% 4	0.9% 1	1.8% 2	0.0% 0	0.0% 0	0.0% 0	100% 113

Significance values are determined using the Holm correction for multiple comparisons. For comparability with traditional partisanship, this table uses θ^{+}, which sets all negative values of θ to zero.

The main message from Table 6 (and Appendix C) is that for about 50% of the cases in which the traditional PID measure identifies a respondent as an independent, we do so as well. Further, most of the remaining traditional independents only have a single-party attachment by our measure. Likewise, for traditional single-party partisans, we identify only 20% as independents and another 27% as single-party partisans. That is, 53% of Danish traditional, single-party partisans have multiple partisan identities in our data.

[53] See Appendix F.26.

4.3 Patterns of co-partisanship

While the analysis above gives us a good sense of the overall level of multi-party partisanship, it does not identify which parties attract co-partisans or account for differences in the strength of co-partisanship for different party pairs. In this section, we use the definition of co-partisanship explained in Section 3.3.3 to develop a measure of the strength of co-partisanship, so that we can explore these questions. (Appendix C provides a similar analysis of total levels of partisanship – including both single-party and co-partisanship.)

As previewed in Section 3.3.3, we can think of each respondent's partisanship as a vector in J-dimensional space (where J is the number of parties in the system). As such, we can also define, for each respondent, all the 2-dimensional sub-vectors in this space, each of which provides a simple geometric interpretation of what it means for a respondent to be "co-partisan" with respect to the corresponding party-dyad (where we only consider positive partisanship for this purpose): a respondent is strongly co-partisan with respect to parties *j* and *k* when she is both equally attached to *j* and *k* and strongly so. As her partisan attachments to either *j* or *k* become less strong or unequal, we should think of her as a less strong co-partisan of *j* and *k*.

Geometrically, we can represent a strong co-partisan as a 2-component vector at a 45-degree angle from the origin, with a large magnitude (given the metric on which θ_{ij}^{+} is estimated, the maximum such magnitudes are around 3).[54] Our measure of co-partisanship strength is the magnitude of this vector times a number that moves continuously between 0 and 1 depending on how closely the vector's direction is to 45-degrees.[55, 56]

In the next sections we explore patterns in this measure, beginning by examining its relationship with voters' left-right images of the parties.

4.3.1 Left-Right and multi-party partisanship

We argued in Chapter 2 that multiple party attachments should be strongly organized by the left-right in multi-party systems. This is because the relative positions of parties on the left-right dimension effectively summarizes the extent to which any two parties are similar on a set of variables that drive the development of multi-party partisan identities, following Lee, Santoso and Stevenson (2024): past patterns of partisan cooperation and conflict (i.e., pairs of parties that regularly cooperate will tend to have more and stronger positive co-partisans); the policy positions and broad values of the parties (i.e., pairs of parties with shared policy goals and values will tend to have more and stronger positive co-partisans), and the pattern of social group support for the parties (i.e., pairs of parties that's draw their support from the same social groups in society will have more and stronger positive co-partisans). Here we explore those hypotheses using our

[54] See Appendix F.27.
[55] The formula we use is $1 - (abs(\lambda_{i(jk)} - 45)/45)$.
[56] See Appendix F.28.

measure of co-partisanship (full results in Appendix D.2). Appendix B provides results from several other methods to explore the structure in multi-party partisanship, including unsupervised clustering methods, which we used to look for natural groupings in the estimated attachment vectors across respondents. As it turns out, these methods tell much the same story that we report in this section – that the left-right images of parties strongly structure patterns of multi-party attachments across parties, providing a robustness check for the results reported here.

Figure 13 provides our first cut at understanding the extent to which voters' left-right images of parties structure patterns of co-partisanship. The data at the respondent-party dyad level. The x-axis in the figure is the respondent's perceived left-right distance between the parties in the dyad, and the y-axis is our measure of the respondent's strength of co-partisanship for the party dyad.

[Figure 13 here]

The figure clearly shows a strong negative relationship between perceived left-right distance and co-partisanship in each country. That said, these patterns are less indicative of a simple linear relationship than they are of a necessary, but not sufficient, condition: when respondents perceive the left-right distance between two parties to be large, they almost always have negligible or weak levels of co-partisanship for those pairs of parties; however, when they perceive two parties to be close together, co-partisanship for those parties may or may not be strong. This is exactly what we should expect, as Figure 13 includes all party pairs for all respondents without accounting for voters' own left-right placement. Consequently, for many party pairs, respondents will report the parties as having similar left-right positions but will have little attachment to either of them (e.g., a leftist voter responding about two rightist parties), contributing many zeros for co-partisanship strengths for ideologically close party pairs depicted in Figure 13.

Another way to examine these data, which de-emphasizes the many zeros included for each party pair, is to aggregate the co-partisanship measure (over respondents) by party dyad. Figure 14 does this by plotting the mean level of co-partisanship of each party pair against their average (perceived) left-right distance from each other. While Figure 13 already indicated that voters will develop more positive co-partisanship for party pairs that are closer together on the left-right, it is important to recognize than not all "distances" between parties on the left-right are necessarily the same. Most importantly, we can ask if, for a given distance between parties, this span crosses the ideological center. While in a pure spatial account of what the left-right means to voters, this would not matter, more recent scholarship (e.g., Nicholson et al. 2018) points to the hypothesis that co-partisanship for parties that cross the ideological center should be depressed relative to other party pairs, separated by a similar distance on the left-right, that do not. In Figure 14, we highlight this: red pairs indicate both parties are on the left, blue pairs are both on the right, and green pairs have one party from the left and one from the right.

[Figure 14 here]

This figure not only highlights the strong negative association between differences in the left-right images of parties and levels of co-partisanship, but also reveals that:

(1) High levels of co-partisanship are mostly restricted to pairs of parties in the same broad ideological family.

(2) While we can discern some asymmetry in levels of co-partisanship for left and right parties respectively, the extent and direction of this asymmetry varies by country. Specifically, in Italy co-partisanship is much higher on the right (and among parties perceived as closer together on the left-right). This pattern is reversed in the Netherlands, while Sweden and the UK show little asymmetry.

(3) Average levels of co-partisanship "bottom out" for most of our countries when parties are separated by about 4 points on the left-right scale. This accounts for the flattening of the loess line on the bottom right in each graph (except Italy) and suggests that positive levels of co-partisanship for party pairs beyond that is probably just noise.

While Figures 13 and 14 provide a high-level look at co-partisanship, "zooming in" on one country reveals other interesting patterns to which we may want to devote more attention. For example, Table 7 examines co-partisanship in Sweden (others in Appendix C), but rather than focusing on our measure of co-partisanship strength, it provides the proportion of respondents with statistically significant θ_{ij}^{+}'s for both parties in each party pair.

Table 7: Proportion of positive co-partisans for pairs of parties in Sweden

	Left Party (1.89)	**Greens** (3.32)	**Soc. Dem** (3.63)	**Center** (4.79)	**Liberals** (5.54)	**Ch. Dem** (6.47)	**Moderate** (7.42)	**Sw. Dems** (7.58)
Left Party		18.6%	12.9%	11.9%	**7.0%**	**5.3%**	**0.9%**	**5.4%**
Greens			11.5%	13.8%	**8.2%**	**4.9%**	**1.2%**	**4.8%**
Soc. Dem				7.1%	**2.8%**	**1.7%**	**0.9%**	**2.0%**
Center					**10.7%**	**7.2%**	**3.0%**	**4.6%**
Liberals						11.2%	5.7%	8.9%
Ch. Dem							7.1%	15.9%
Moderate								8.5%

Cell entries are the proportion of respondents who had a statistically significant theta (applying a Holm correction for multiple comparisons) for both the row and column party. The bolded numbers are for party dyads that cross the center of the left-right scale (based on average perceived party placements, which are the numbers in the second row of the column titles). This table uses θ^{+}, which sets all negative values of θ to zero.

The estimates in Table 7 help quantify the extent of co-partisanship across the center of the left-right. As expected, this kind of co-partisanship is rare -- only about 4.4% overall (4.0% if we exclude the Center-Liberal dyad); while co-partisanship among the parties within the broad left and right ideological families are much more common.

Within the left and right families, other patterns are also apparent. Examining co-partisanship across the left-right divide (the bolded numbers), the Moderates and Social Democrats stands out with notably low rates: only 1.5% and 1.9% for the Moderates and Social Democrats, respectively, compared to the average of 5.8% for other parties. Indeed, this dyad had the lowest co-partisanship – 0.9% – among all party pairs, even though they are hardly the farthest apart on the left-right. Given their traditional roles as the leading parties of the right and left blocs, this may suggest that the development and maintenance of cross left-right co-partisanship is particularly difficult to achieve when one of the parties in question leads the opposite bloc.[57]

In Appendix C, we report several analyses that do not rely on our co-partisanship strength measure but simply ask the extent to which the θ_{ij}'s covary with one another. These include confirmatory multi-dimensional IRT models (MIRT) to estimate correlations among θ_{ij}'s for different parties and unsupervised clustering methods to find natural groupings among our respondent's the θ_i vectors. This clustering analysis reveals a set of common voter "types," including large groups of individuals who are broadly attached to parties of the left or the right – as we have also seen in the results above – but within these broad categories, we also see distinct sub-types – those who are more strongly attached to ideologically extreme parties and those who are more strongly attached to the ideologically central parties (within each ideological family). We also find a set of "true" centrists (who are most attached to centrist parties, with attachment falling off as parties get more extreme on the left or right) and independents (who are not attached or only weakly attached to any party).

4.3.2 Why Left-Right Ideology: unpacking the influence of the left-right on co-partisanship

It is clear from the analyses above that multi-party PSI is strongly structured by parties' relative left-right positions. In our theoretical discussion, we suggested that this is expected because parties' relative left-right ideologies summarize (in the minds of voters) corresponding similarities and differences in policy priorities, broader values, social bases of support, and patterns of partisan conflict and cooperation (i.e. which parties are allies and which are enemies), and it is similarity in these factors across parties that we think leads to the development and maintenance of multi-party identifications. With our estimates of PSI, we can now begin to explore this directly.

Specifically, we have data on one important component of this story for two of our countries: the extent to which voters perceive party pairs as cooperating. Cooperation between parties is likely the primary means through which voters develop co-partisan attachments, so it should be strongly associated with levels of co-partisanship. Indeed, to the extent that there is some slippage in the impact of this variable on voter's left-right images of parties, we might even expect this association to survive a control for the perceived left-right distance between the parties (an expectation we test in Table 8 below).

We draw on two data sources. First, for Denmark we combine our θ_{ij} estimates with Lee, Santoso, and Stevenson's (2025) data on perceived cooperation between party pairs in 2019. In this data,

[57] See Appendix F.29.

respondents rated all the different pairs of Danish parties on a 0 ("Never cooperate") to 10 ("Almost always cooperate") scale, in response to the question "How often do you think these two parties cooperate with each other in Denmark's federal politics?" Since these data come from a different survey than the Danish survey we have been using so far (fielded at the end of 2021), we can only correlate co-partisanship with perceived cooperation at the aggregate level (i.e., averaged over respondents for each party-dyad).

The second data source is a survey we conducted in the UK in 2024. That survey included the same party cooperation question as in the Danish Survey as well as out multi-party partisanship scales, allowing us to examine the relationship between perceived partisan cooperation and co-partisanship at the individual level.

Figure 15 graphs the aggregated Danish data and Figure 16 the individual-level UK data. Both figures confirm a strong bivariate association between perceived levels of cooperation between parties and respondents' strength of co-partisanship for the pair.

[Figure 15 here]

[Figure 16 here]

While expected, this result was not inevitable. That is, one can readily imagine the opposite: that experiences of governing together might reduce co-partisanship. Why might we expect this? Governing in cabinets, as the large literature on cabinet governance attests (e.g., Martin and Vanberg 2011), often involves difficult policy compromises, attempts by cabinet parties to "renege" on coalition bargains, counter efforts of cabinet partners to monitor and stop such efforts, and in many cases, conflictual cabinet dissolutions. Thus, it is not obvious that the experience of governing together in cabinets would necessarily create positive feelings among cabinet partners and, by extension, their voters. However, to the extent that such conflicts are hidden from most voters (who are likely not paying enough attention to follow all the nuances of intra-cabinet conflicts), the more symbolic unity of just being on the same team may be enough to drive positive affective reactions to this kind of cooperation.

We partially extend the above analysis to the other countries in our sample using historical cabinet participation data to answer: do party dyads who have historically cooperated in cabinets have more co-partisans? To answer this, we calculated the percentage of months since January 1980 each pair of parties had governed together in Denmark, the Netherlands, Italy, and Sweden. Figure 17 presents this data against levels of co-partisanship, confirming the positive association between partisan cooperation and co-partisanship. Since this analysis does not measure cooperation with a perceptual variable asked in the same survey as our partisanship batteries, the variable on the x-axis – the number of months governing together – is exogenous (and pre-treatment) with respect to individual attachments to parties. This reassures us that the association between partisan cooperation and co-partisanship that we found in the survey data is not entirely spurious.

[Figure 17 here]

That said, the fit between cabinet cooperation and co-partisanship is looser than with the perceptual data. This could be due to the endogeneity in the perceptions but may also stem from the fact that a lot of cooperation between parties happens outside formal cabinet roles, requiring data on informal cooperation that is not based on respondent perceptions.[58]

4.3.3. Summary and Integration

We end this section with a multivariate look at the three predictors of co-partisanship highlighted above and in our theoretical discussion: partisan cooperation, shared social group support, and shared policy positions and values. We take advantage of original data from two Danish Surveys that allow us to measure voters' perceptions of each of these variables, as well as levels of co-partisanship, for all the Danish party dyads.

The first of these is our survey fielded in 2021, from which we calculate average levels of co-partisanship (across respondents) for each of the 45 Danish party dyads. The second survey, fielded by several of the authors in 2019, did not include our multi-party attachment batteries but asked respondents about their perceptions of parties' policy positions on eight prominent issues, five broader values dimensions, perceived cooperation between party pairs, and support from sixteen different social groups. Using these data, we constructed summary measures of the extent to which party dyads in Denmark were dissimilar on each dimension. For policy/values, for instance, we calculated the perceived Euclidean distances between parties in the pair in the 13-dimensional policy space implied by our 13 policy/values questions. Appendix D.2 provides full details and demonstrates that these measures produce sensible estimates of the voters' perceived levels of cooperation, social-group similarity, and policy and values similarity for each party dyad.

Combining the 2019 data with our 2021 co-partisanship measures of average co-partisanship for each party dyad in Denmark allows us to examine, in a multivariate model, how co-partisanship responds to voters' perceptions of (1) conflict and cooperation between the parties, (2) the similarity of their policy preferences and values, and (3) the similarity of the social groups from which they draw their support. Table 8 presents estimates with and without party fixed-effects, with variables transformed to comparable scales (see Appendix D.2).

[58] See Appendix F.30.

Table 8: The conditional association of cooperation, perceived policy/values similarity, and perceived similarity to social group support and levels of co-partisanship in Denmark

	(1)	(2)
Perceived level of cooperation between parties in dyad	0.19*	0.22*
	(0.03)	(0.04)
Distance between the perceived patterns of social group support between parties in the dyad (16 social groups)	–0.09*	–0.04
	(0.02)	(0.03)
Distance between the perceived policy positions and values of parties in the dyad (8 policies and 5 values)	–0.04	–0.06*
	(0.03)	(0.02)
Alternative		0.01
		(0.01)
Conservative People's Party		–0.003
		(0.01)
Danish People's Party		0.01
		(0.01)
Green Left		0.004
		(0.01)
New Right		0.01
		(0.01)
Social Democrats		–0.03*
		(0.01)
Social Liberal Party		–0.01
		(0.01)
Venstre		–0.01
		(0.01)
Unity List		0.02*
		(0.01)
Intercept	0.09*	0.06
	(0.03)	(0.04)
Observations	44	44
Adjusted R^2	0.91	0.95

The results of this analysis largely confirm that each of the three hypothesized conditional associations is readily apparent in these data. Partisan cooperation is the largest and most robust of these associations, consistent with our intuition that cooperation (in and out of government) is the primary path by which co-partisanship is created and maintained in multi-party democracies. Both of the other variables in the model are in the predicted direction – greater dissimilarity in policy/ values stances or social groups support is associated with depressed levels of co-partisanship – though the estimates fall just short of statistical significance in one or the other specifications.

5. Evidence for the usefulness of multiparty identity

We now assess whether our multi-party partisanship measure enhances our understanding of the political behavior of citizens in multi-party democracies. Specifically, we ask whether it helps explain (or is otherwise associated with) political behavior that traditional measures of party identification can say little about -- including that of complete independents and, for some attitudes and behaviors, non-partisans with respect to a given party.

5.1 Does $\vec{\theta}_i$ predict political attitudes and behaviors among traditional independents and non-partisans with respect to a given party?

We examine the group of respondents who traditional measures classify as independents (i.e. they report no party to which they feel closest). As we saw in Table 5, about a quarter of our respondents are traditional independents. These respondents are particularly probative because traditional measures assign them a 0 for all parties, yet our measure often reveals non-zero (and relatively large) θ_{ij}'s for at least some parties.

Traditional independents, however, are only one kind of useful comparison we can make between traditional measures and ours. Given the construction of the traditional measure, in an 8-party system like Sweden, for example, respondents will have either 7 or 8 zeros in the traditional measure: 8 zeros for complete (traditional) independents and 7 zeros for partisans who are (by definition) non-partisan with respect to all the parties but one. Thus, traditional PID assigns a zero to many respondent-party pairs, while we produce a potentially non-zero estimate for each respondent-party pair.

This presents the opportunity to do an analysis that other studies of single-party, multi-item scales of PSI have not been able to do: we can examine whether variation in our PSI measure across respondents (and sometimes respondent-parties) predicts outcomes like vote choice, political engagement, political efficacy, and party like/dislike ratings, both among traditional independents (i.e., all respondent-party pairs are zero) and among parties to which traditional partisans are not attached (i.e., where the traditional measure is zero for that subset of respondent-party pairs).

If our measure is predictive of a wide variety of political behaviors among respondents with the same (zero) coding on the traditional measure, we should be persuaded that multi-party PSI captures something the traditional measures miss. In contrast, if these relationships are weak and insignificant, then we should conclude that no-zero estimates multi-party attachments for these traditional non-partisans are just noise.

5.1.1 Vote choice

Our first test estimates a conditional logit model of vote choice over parties, examining whether our θ_{ij}'s improve predictions of vote choice among traditional independents. Table 9 reports the results.

Table 9: Multi-party ID and vote choice among traditional independents

	Denmark	Italy	Netherlands	Sweden	UK
θ_{ij}	2.00*	1.61*	2.07*	1.88*	2.16*
	(0.32)	(0.25)	(0.30)	(0.32)	(0.27)
# of Respondents	59	86	65	54	111
Pseudo R^2	0.24	0.19	0.29	0.27	0.26

The observations are traditional independents who cast a vote for one of our parties (e.g., in Denmark we have 113 traditional independents in the sample, of which 59 reported the intention to vote for one of the parties for which we estimated thetas. The model is a conditional logit in which the various parties are the choice set. Numbers in parentheses are standard errors.

Clearly, among these traditional independents, our measures of multi-party attachment are statistically significant and substantively large predictors of vote choice in all five countries. Figure 18 uses the coefficients from Table 9 to graph the predicted probability of voting for a party across the full range of θ_{ij}. The results show a very strong association between θ_{ij} and the probability of voting for that party among traditional independents. Likewise, the pseudo R-squared is fairly large, which is notable given that traditional independents are typically less politically engaged in the first place.[59]

These findings mirror those presented by Garry (2007: 354), who used a unique measure of multi-party PID and top conclude (for the Northern Irish case) that, "these new party identification measures are useful predictors of voting, adding information over and above that provided by the conventional measure."

[Figure 18 here]

5.1.2 Party feeling thermometers

Our next test, reported in Table 10, examines the impact of θ_{ij} on traditionally independent respondents' like/dislike ratings of the parties (feeling thermometers). Once again, we find a very strong relationship in each country, with substantial variance explained.

[59] This analysis is conditional on these traditional independents reporting who they would vote for if the election were held today. As expected, many said they would not vote (so the Ns are smaller than the full sample of traditional independents used elswhere).

Table 10: Multi-party ID and party thermometer (like/dislike) scores for parties among traditional independents

	Denmark	Italy	Netherlands	Sweden	UK
θ_{ij}	21.84*	23.88*	21.55*	21.05*	19.58*
	(1.07)	(0.93)	(0.94)	(1.22)	(1.26)
Intercept	36.62*	30.82*	37.76*	35.50*	36.95*
	(1.74)	(1.19)	(1.28)	(1.58)	(1.23)
Observations	1,017	1,329	1,341	776	1,141
# of Respondents	113	223	149	97	163
Adjusted R^2	0.26	0.34	0.30	0.34	0.21

The observations are all respondent-party dyads for traditional independents. The model is a linear regression. Numbers in parentheses are individual-level clustered standard errors. Stars indicate the coefficient is significant at p < .05.

In addition to independents, we can examine the correlation between θ_{ij} and the thermometer scores for the parties to which traditional partisans are (according to that measure) not attached. Using the data in respondent-party format, we drop, for each traditional partisan, the one party for which the traditional PID measure is not 0. Thus, the analysis below only includes respondent-party pairs for which traditional PID indicates the respondent is not a partisan of that party.

The results in Table 11 are striking: θ_{ij} accounts for 40% of the variance in traditional partisans' thermometer ratings for parties they did not report being attached to. This clearly demonstrates that our PSI measures add substantial empirical value beyond traditional measures.

Table 11: Multi-party ID and party thermometer (like/dislike) scores among traditional partisans (but excluding the party to which they are traditionally attached)

	Denmark	Italy	Netherlands	Sweden	UK
θ_{ij}	22.74*	23.82*	21.26*	21.79*	20.98*
	(0.35)	(0.41)	(0.35)	(0.34)	(0.48)
Intercept	37.49*	33.74*	41.11*	34.36*	37.09*
	(0.52)	(0.53)	(0.46)	(0.41)	(0.50)
Observations	7,939	4,709	7,020	6,411	6,109
# of Respondents	982	917	858	910	1015
Adjusted R^2	0.37	0.43	0.39	0.43	0.35

The observations are all respondent-party dyads for traditional partisans, excluding the party for which the traditional PID measure is non-zero. The model is a linear regression. Numbers in parentheses are individual-level clustered standard errors. Stars indicate the coefficient is significant at $p < .05$.

5.1.3 Political participation and engagement

A great deal of work has shown that identifying more strongly with a (single) party "encourage[s] a person to become active in the political process to support his or her side." (Dalton 2016: 7; Jung 2017; Moral 2017). For example, using a multi-item battery to measure single-party partisan identity, Huddy, Mason, and Aarøe's (2015) show that individuals with stronger PSI are more involved in political campaigns and more emotionally reactive to campaign-related events.

In our survey, we included a nine-item political participation battery, from which we built an IRT-model based scale that covers the underlying trait well (see Appendix D.4).

Our analysis differs from previous ones in two ways. First, since participation is a respondent characteristic, we examine its association with θ_{ij} and participation at the respondent level (for traditional independents) rather than at the respondent-party level. This means making some choice about how the J estimated θ_{ij}'s for each respondent will enter the respondent-level model. As discussed in Section 3.2.5, the natural choice is $||\vec{\boldsymbol{\theta}}_i||$ (and is positive or negative variants), which is the length of the respondent's partisan attachment vector.

Second, unlike most studies that have examined the association of partisan strength with participation, we focus on examining the relationship among traditional independents. Asking if $\vec{\boldsymbol{\theta}}_i$ predicts participation among traditional independents is a particularly high bar since independents are known to participate less – they are much less likely to vote, work for campaigns, and donate to candidates. We therefore do not expect to be able to explain much of the variance in this sample and instead focus mainly on whether we estimate statistically significant associations between participation and partisan identity strength (positive, negative, and overall).[60] Table 12 provides the results.

Table 12: Multi-party ID and political engagement/participation among traditional independents

	Denmark		Italy		Netherlands		Sweden		UK	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
$\|\|\overrightarrow{\theta_i^+}\|\|$	-0.08		0.01		0.02		0.09		0.18	

[60] The same observation is less applicable to the voting models in the previous sections, which had large pseudo R-squared estimates, because those estimates are conditional on the respondent intending to vote for one of the parties.

	Denmark		Italy		Netherlands		Sweden		UK	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
$\|\|\overrightarrow{\theta_i^-}\|\|$	(0.10)		(0.07)		(0.07)		(0.08)		(0.10)	
	0.27*		0.14*		0.13*		0.07		-0.10	
	(0.06)		(0.05)		(0.04)		(0.06)		(0.07)	
$\|\|\overrightarrow{\theta_i}\|\|$		0.22*		0.15*		0.12*		0.09		0.002
		(0.06)		(0.04)		(0.05)		(0.06)		(0.07)
Intercept	-0.23*	-0.36*	-0.14	-0.20*	-0.17	-0.20*	-0.18	-0.18	-0.05	-0.003
	(0.12)	(0.12)	(0.09)	(0.08)	(0.10)	(0.10)	(0.14)	(0.14)	(0.11)	(0.11)
Observations	113	113	223	223	149	149	97	97	163	163
Adjusted R^2	0.15	0.11	0.03	0.04	0.04	0.04	0.003	0.01	0.02	0.00

Respondents are traditional independents. The model is linear regression. Numbers in parentheses are standard errors. Stars indicate the coefficient is significant at $p < .05$.

Table 12 shows the expected directional pattern for positive partisanship (i.e., stronger partisanship leads to more participation), but the estimates are not statistically different from zero in all countries. This pattern is unsurprising given the overall lack of engagement among independents. What is more surprising is the strong negative association between increasing negative PSI and political participation. This contradicts recent speculation in the small literature on negative partisanship that negative partisanship will depress participation and engagement. Instead, these findings are consistent with Bankert's (2024) result that both positive and negative (single-party) partisanship increases political participation. This result becomes particularly interesting when compared to the results on political efficacy below.

5.1.4 Political efficacy

All our surveys included a question on respondents' political efficacy, asking "How much would you say that the political system in [COUNTRY] allows people like you to have a say in what the government does?" (with answer categories ranged from "Not at all" to "A great deal"). Table 13 provides the estimated associations between a respondent's total positive and negative partisan attachment and political efficacy (from an ordered logistic model in which the dependent variable was the 5-category response to the question above).

Table 13: Multi-party ID and political efficacy among traditional independents

	Denmark	Italy	Netherlands	Sweden	UK
$\|\|\overrightarrow{\theta_i^+}\|\|$	0.27	0.77*	0.32	0.58*	0.57*
	(0.24)	(0.19)	(0.19)	(0.28)	(0.24)
$\|\|\overrightarrow{\theta_i^-}\|\|$	-0.22	-0.59*	-0.44*	-0.45*	-0.28
	(0.15)	(0.13)	(0.14)	(0.18)	(0.15)
1|2	-1.16	-0.32	-0.96	-1.23*	-1.14*
	(0.32)	(0.22)	(0.27)	(0.41)	(0.29)
2|3	-0.07	1.19*	0.58*	0.30	0.37
	(0.30)	(0.24)	(0.27)	(0.38)	(0.27)
3|4	1.71*	2.66*	2.97*	2.29*	1.93*
	(0.34)	(0.31)	(0.44)	(0.47)	(0.32)
4|5	2.76*	4.36*	3.84*	3.74*	2.99*
	(0.45)	(0.57)	(0.62)	(0.69)	(0.41)
Observations	113	223	149	97	163

Observations are traditional independents. The model is an ordered logit. Numbers in parentheses are standard errors. The numbers labelled, for example, "1|2" are the estimated intercepts in the ordered logit. Stars indicate the coefficient is significant at p < .05.

These results are interesting and novel. While Table 12 showed that as traditional independents become more negatively (or insignificantly positively) attached, they report greater participation, this result in Table 13 reveals that feelings of political efficacy (among traditional independents) are positively associated with positive PSI but negatively with negative PSI. This adds some subtlety to the idea in the literature that links negative partisanship to disengagement and disillusionment with politics. Specifically, it suggests that the hypothesized negative impact of negative partisanship may to limited to feelings of efficacy (or disillusionment) around having a positive impact on what government does, while encouraging (a disillusioned form of) engagement – e.g., working to stop your enemy rather than to achieve something positive.

5.1.5 Summary

The message of this section is clear. If θ_{ij} captures PSI, then these identities are associated with a variety of political behaviors and attitudes, even when we focus on the hardest case – independents and non-partisans with respect to a given party. The strong, significant associations between our θ_{ij} and a variety of political attitudes and behaviors (and the that our θ_{ij} account for a large percentage

of the variance in some of these behaviors) strongly support the idea that multi-party PSI adds considerable value to traditional, maximal, and single-party measures of PID.

5.2 Does θ_{ij} predict levels of motivated reasoning?

In the previous sections, we examined how θ_{ij} are structured and whether they could explain the political attitudes and behaviors of traditional independents and non-partisans (with respect to a specific party). We now begin a broader comparison of our measure to traditional single-party PID, focusing on the extent each is consistent with the "calling-card" of PSI: partisan motivated reasoning.

Motivated reasoning is a broad term that covers a variety of ways individuals who identify with a group alter their beliefs (and information processing more broadly) to favor it. This can include "wishful thinking" or "defensive reasoning" about party performance, i.e., partisans believe (or at least report) that their party has performed better than it really has. Indeed, motivated reasoning is so central to what it means to have a partisan social identity that it was a key part of Campbell et al.'s (1960, p. 133) original statement of the idea:

> *"Identification with a party raises a perceptual screen through which the individual tends to see what is favorable to his partisan orientation. The stronger the party bond, the more exaggerated the process of selection and perceptual distortion will be."*

Dramatic examples of motivated reasoning are easy to find, including studies that show large biases in partisan's views of retrospective economic performance (Bartels 2002; Gerber and Huber 2010; Lavine, Johnston, and Steenbergen 2012; Santoso 2020; Tilley and Hobolt 2011), expectations about future performance (Duch and Stevenson 2011), evaluations of politicians (Goren 2002; Lebo and Cassino 2007), beliefs about a variety of other politically relevant facts (Bartels 2002; Bullock et al. 2015), memory for partisan facts (Jerit and Barabas 2012), and how evidence is evaluated (Kahan et al. 2017).

Motivated reasoning is also readily observable in partisans' expectations about what is likely to happen in elections and cabinet formation. Huddy, Bankert, and Davies (2018), for example, used a (single-party) PSI battery to show that, just before the 2015 UK election, strong partisans of the Conservative and Labour parties were much more likely than weak partisans to believe their party would win the upcoming election. Further, Huddy, Davies, and Sandor (2020) report on two analyses of Swedish and Dutch voters in which their expectations about which parties were likely to obtain positions in cabinet following an upcoming election were strongly influenced by their PID.

Left-right self-placement is another variable often impacted by partisan rationalization. The comparative politics literature on the drivers of left-right self-placement has consistently shown that identification with a party impacts voters' left-right self-placement as much as (or more than) policy positions – a finding that is indicative of a kind of partisan rationalization (Inglehart and Klingemann 1976; Knutsen 1995; Lesschaeve 2017; Medina 2015).

In the following sections, we test whether our multi-party measure reveals motivated reasoning among multi-party partisans, including cases where we only include respondent-party pairs in which traditional PID is zero. If we find evidence of motivated reasoning, this should be considered evidence that our measure is capturing PSI among these multiple partisans, as motivated reasoning is a strong marker of social identity.

5.2.1 Partisan rationalization in beliefs about incumbent performance

The most prominent form of partisan rationalization in the literature concerns beliefs about the performance of incumbents, with the usual hypothesis being that stronger partisans of the incumbent party should evaluate incumbent performance on valence issues (like the economy and crime) more positively than less strong partisans or non-partisans (either via direct cognitive biases and/or selective consumption of information). For example, Duch and Stevenson (2010) showed, in an analysis of inflation expectations in ten democracies, that partisans across Europe expect inflation to be lower when the party they support is in power. Likewise, most scholars who study negative partisanship would expect stronger negative partisanship with respect to an incumbent to produce more negative evaluations of incumbent performance.

Our 2024 UK survey included questions probing respondents' retrospective and prospective beliefs about economy (covering growth, inflation, and unemployment) and crime (a particularly prominent issue in the UK at the time of our survey).[61] IRT scaling of these eight items produced a strong unidimensional scale, which we use in the analysis below (see Appendix D.5 for details). Analyses using individual items yield the same conclusions.

Figure 19 confirms the usual result using the traditional PID measure.

[Figure 19 here]

Unlike most previous studies of partisan reasoning about incumbent performance, our measure of identity with the incumbent (Labour) varies continuously across all British respondents, from strong negative attachment to strong positive attachment. Thus, to test if our measure "works" in the same way as maximalist measures of PID across the full range of positive and negative attachments to an incumbent, we need to check if the kind of motivated reasoning effect that we see in Figure 19 is observable over the whole range of θ_{iLAB}.

We examine the relationship between θ_{iLAB} and beliefs about performance non-parametrically. Doing so allows the data to tell us if, for different ranges of θ_{iLAB} , we see the expected positive relationship. For example, if we only discovered motivated reasoning about economic performance for large values of θ_{iLAB} then we would conclude that our measure of attachment to the incumbent Labour party does ***not*** have proportionate effects on evaluations of economic performance across

[61] The Labour government had come to power approximately six months before our survey was in the field, so the time frame which we asked the retrospective economy questions was six months.

all levels of partisan attachment scale, but only for strong incumbent partisans that traditional PID measures likely already identify.

Figure 20 presents a local kernel regression of each respondent's performance score on her attachment to Labour. Local kernel regression allows us to visually assess whether motivated reasoning seems to happen in proportion to the size (and direction) of θ_{iLAB} along its entire range.[62]

[Figure 20 here]

The result clearly supports the idea that our respondents' attachment to Labour has the predicted positive relationship with performance evaluations across the whole range of θ_{iLAB}. Given that one of the clearest (and well documented) effects of partisan identity is motivated reasoning in favor of the party to which one is attached, we think the results in Figure 20 provide strong evidence that our measure of identity strength for the incumbent reflects real variation in the strength of partisan identity across its full range.

With our multi-party measures, we can also ask new questions about the way motivated reasoning about incumbent performance depends on multi-party attachments. For example, do individuals who are not attached strongly to Labour, but who are attached strongly to other leftist parties, appear to engage in motivated reasoning about performance, or does such motivated reasoning require an identity with the incumbent specifically?

Table 14 addresses this via a regression of incumbent performance on respondents' attachments to *each* party. In this analysis, we focus only on respondents who knew (and told us in the survey) that Labour held the prime ministry.

Table 14: The relationship between θ_{ij} and our index of incumbent performance evaluations among respondents who knew that Labour held the prime ministry

	θ_{ij}^{+}	θ_{ij}^{-}
Labour	0.51*	-0.67*
	(0.06)	(0.06)
Lib Dems	0.12	-0.10
	(0.07)	(0.07)
Greens	-0.003	-0.01
	(0.07)	(0.06)
SNP	0.07	-0.04
	(0.08)	(0.07)
Plaid Cymru	0.13	0.13
	(0.10)	(0.08)
Conservative	-0.01	-0.15*
	(0.07)	(0.06)

[62] See Appendix F.31.

	θ_{ij}^{+}	θ_{ij}^{-}
Reform	-0.43*	0.30*
	(0.07)	(0.06)
Intercept	-0.14*	0.28*
	(0.06)	(0.05)
Observations	741	741
Adjusted R^2	0.20	0.25

Coefficients in each column are from a regression of θ_{ij} for each party in the left most column, on our index of beliefs about incumbent (Labour) performance (higher numbers indicate better performance). Stars indicate the coefficient is significant at $p < .05$.

Looking first at Model 1 (which use $\theta_{ij}^{+}{}'s$ as the predictors), the results clearly show that attachments to Labour conditioned evaluations of economic performance over the sixth months Labour had been in office. The coefficient on $\theta_{i\,LAB}^{+}$ is more than triple the size of other positive coefficients for other parties and is the only one that is statistically significant.

This is an important result, not only because it confirms a high degree of motivated reasoning is apparent using our measure of attachment to Labour, but also because *it tells us that attachments to other parties, which may be aligned with Labour in various ways, do not readily "transfer" to motivated reasoning about Labour –* ***it takes attachment to the Labour itself***.

This rules out, at least for the UK, the otherwise reasonable possibility that respondents would treat Labour's incumbency as a kind of "pan-leftist" government, in which the leftist parties have heightened influence (and so some responsibility) for outcomes. If this were the case, however, partisans of other leftist parties should engage in motivated reasoning for the left generally – and thus for Labour when it is in power (as the leading party of the left). The UK data show that this is not the case.[63]

Model 1 also reveals a strong negative association between performance evaluations and $\theta_{i\,REFORM}^{+}$: as a respondent identifies more strongly with the Reform party, the worse he or she evaluates economic and crime performance during Labour's incumbency, exactly what we expect if, as is quite likely, partisans of Reform think of Labour as an important "out-party."

Turning to Model 2, which examines how negative partisanship impacts performance evaluations, we again find that increasing negative attachment to Labour is strongly associated with more negative evaluations of the economy and crime during the Labour's incumbency. Coefficients for negative attachment to other leftist parties are all small and statistically insignificant. Negative attachment to the Reform party shows the reverse pattern – greater negative attachment to Reform is associated with better performance evaluation – , mirroring our result from Model 1.

The one anomalous result in Model 2 is the small negative coefficient on negative attachment to the Conservatives. This result suggests that stronger negative attachment to the Conservatives is

[63] See Appendix F.32.

associated with worse performance evaluations, just as we found with Labour (though it is about ¼ the size of Labour's effect). Indeed, looking back at the Model 1 result for the Conservatives showed no association with performance evaluation, where we might have expected a strong negative coefficient. This suggests that some respondents treated the Conservatives less like an opposition party (and more like a government party) than we expected.

Part of what is going on here may have to do with the fact that when our survey was in the field, the Labour government had been in power only about 6 months, following an extended period of Conservative rule. As Fortunato and Stevenson (2013) demonstrated for the Netherlands, when long-time patterns of government composition change, there can be a long-lasting hangover effect in which some voters, paying limited attention to politics, continue to treat the party of the previous PM as if it were still in power. For example, Fortunato and Stevenson demonstrated that the economic vote in the Netherlands in 1998 was depressed by almost a third due to about 25% voters not knowing that the long line of Christian Democratic prime misters had come to an end.[64]

Overall, these results reveal that performance evaluations depend much more strongly on attachment (positively or negatively) to the incumbent Labour party than to other ideologically-compatible parties. Thus, this kind of partisan rationalization appears to require partisan *attachment* rather than ideological similarity. To the extent that attachment to other parties matters at all to these evaluations, it is only attachment (positive or negative) to one of Labour's clearest out-parties: the Reform party.

5.2.2 Partisan rationalization in left-right self-placements

Another commonly observed type of partisan rationalization is the tendency of partisans to adopt the policy positions of their party. While some work has used multi-dimensional policy scales to examine this (e.g., Lesschaeve 2017), the most common empirical approach has examined left-right self-placement to determine whether it is driven more by the respondent's reported policy preferences or partisanship – with the latter indicating partisan rationalization.

While there is some debate in the empirical literature about the relative importance of these drivers of left-right self-placements (e.g., Inglehart and Klingemann 1976; Knutsen 1995; Lesschaeve 2017; Medina 2015), a fair reading of the existing work suggests that partisanship is likely a bigger driver of left-right placements than policy. For example, based on an analysis of left-right self-placement in a large sample of western and eastern European countries, Medina (2015: 775) concluded that "the partisan component" has a stronger effect than other factors, "resulting in further proof of the main role that parties play in influencing voters' positions on the left–right scale."

We explore this by examining the extent to which similarity in policy positions and partisan identities of any two respondents predicts similarity in their left-right self-placements. Specifically, for each respondent, we have J θ_{ij}'s, and J traditional PID scores, where J is the number of parties.

[64] See Appendix F.33.

Thus, each respondent's traditional PID and partisan identity ($\vec{\theta}_i$) can be represented by points in J-space and the similarity of any two voters' partisan identities can be characterized as the Euclidean distance between their party identity vectors in the J-dimensional partisan identity space (for both traditional PID and $\vec{\theta}_i$). Likewise, respondents provided their positions on eight policy issues and so we can represent a combined policy position for each as a point in an 8-dimensional policy space and compute the Euclidian distances between any two respondents in this space.

With this, we first reorganize our data into respondent dyads and compute the distance between the dyad members' respective (1) PSI vectors, (2) traditional PID profiles, (3) left-right self-placements, and (4) policy positions (over 8 policies). We then regress the distance in self-placements on all the other variables, expecting to see a positive impact of the partisan attachment variables, controlling for similarity in policy preference – a result we interpret as evidence of partisan rationalization.

Table 15: Do respondents who have similar $\vec{\theta}_i$'s place themselves closer together on the left-right, controlling for the extent to which they have the same policy preferences?

	Denmark	Italy	Netherlands	Sweden	UK
Distance between J-dimensional vectors $\vec{\theta}_i$ and $\vec{\theta}_{i'}$	0.13*	0.17*	0.12*	0.14*	0.13*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Distance between 8-dimensional policy positions of Rs *i* and *i'*	0.07*	0.05*	0.05*	0.07*	0.07*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Distance between J-dimensional traditional PID vector of Rs	0.03*	0.04*	0.03*	0.06*	0.01*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Intercept	0.37*	0.36*	0.30*	0.09*	0.22*
	(0.01)	(0.004)	(0.004)	(0.01)	(0.003)
# of Respondent Dyads	590,241	621,055	501,501	500,500	693,253
Adjusted R^2	0.08	0.10	0.08	0.13	0.08

The results in Table 15 confirm that respondents with similar policy positions will nevertheless place themselves increasingly farther apart on the left-right dimension as their partisan identities (as captured by $\vec{\theta}_i$ and $\vec{\theta}_{i'}$) diverge. Likewise, two respondents who report quite different policy positions will nevertheless place themselves closer together on the left-right if they share a similar partisan identity.

Of course, there are other ways to examine this kind of motivated reasoning, including using our respondent's left-right placements of parties and themselves. For example, Nasr (2021) used party

like/dislike scales to show that, on average, voters perceive themselves as ideologically closer (on the left-right scale) to parties that they like and more distant from parties they dislike – a result he interpreted as indicative of a kind of motivated reasoning. In Appendix D, we easily replicate this finding replacing like/dislike ratings of the parties with our estimates of $\vec{\boldsymbol{\theta}}_i$.

5.2.3 Partisan rationalization about the outcomes of elections

Another common way in which researchers have investigated motivated reasoning is in expectation formation. For example, Huddy, Bankert, and Davies (2018) have shown, using a single-party, maximalist measure of PID, that increasingly strong partisans expect their party to do increasingly well in elections. If partisans engage in this kind of "wishful thinking" about what *will* happen in elections, an even stronger (and so rarer) form of partisan rationalization concerns what actually did happen. Specifically, if this kind of partisan rationalization occurs, we should expect to see a positive association between our measure of a respondent's PSI for a given party and her beliefs about how well that party did in the last election.

Usefully, in our 2024 survey in the UK we asked each respondent, "After the general election that occurred on July 4th, 2024, what was each party's vote share?" Respondents reported actual percentages for each party, and their answers were constrained to sum to 100% (see Lee, Haime, and Stevenson 2017 for evidence that most European voters have little trouble doing this kind of task and that large numbers are quite accurate). Figure 21 is the result of a regression of their answers to this question on θ_{ij} for each party.[65]

[Figure 21 here]

There is a strong relationship between strength of attachment to each party (both positive and negative) and the extent to which respondents rationalize the election result in favor of (or against) that party. Only for Plaid Cymru (for which we had little attachment in the sample) was the effect not statistically different from zero. This again confirms the idea that our measures of attachment to each party produce the kinds of motivated reasoning predicted by theories of partisanship as a social identity, with its attendant in-group biases. It's also worth noting that when we estimate θ_{ij} to be zero, we get no bias in retrospections about electoral performance, which is compatible with Lee, Haime, and Stevenson's (2017) finding that European voters, on average, have accurate expectations about the outcome of elections.

[65] Full results are in Appendix D.3.

5.2.4 Partisan rationalization in beliefs about which parties are in cabinet

Our surveys in Sweden, Denmark, Italy, and the UK asked each respondent to identify whether each party was in the cabinet or not at the time of the survey.[66] If partisans engage in partisan rationalization with respect to this question, we should expect to see a positive association between our measure of PSI for non-cabinet parties and the likelihood our respondents will think that party is in cabinet.

Compared to studies of voter's expectations or retrospections about vote shares (like those discussed in the previous section), this application of partisan rationalization is an even harder case. First, it does not ask the respondent what she *expects* to be true in the future, but instead what is true now. Likewise, unlike seat shares, information about the identity of the cabinet parties is widely and constantly available in the media, not just following elections but on an on-going basis. Thus, for this kind of partisan rationalization to occur, respondents positively attached to a party not in the cabinet must report something that is both false and quite well known to be so. This is a high bar for mistaken in-group beliefs, but the few studies that ask voters who is in the cabinet have previously found such effects (Lin, Santoso, and Stevenson, 2024).

That said, the characteristics of the cabinets in power at the time of our surveys, and, in one case, the recency of the latest government formation, may make mistaken beliefs about cabinet membership more common in our samples. Specifically, in Sweden our survey was in the field only a month after a new cabinet formation (and the British only 6 months).[67] Further, both the Danish and Swedish cabinets were minority cabinets and so depended on non-cabinet parties for parliamentary support. Finally, the Italian 5-party cabinet was led by an independent PM. Thus, both the relative novelty of some of these cabinets and their arguably ambiguous partisan characters (at least from the point of view of a typical citizen) likely results in more mistaken beliefs about which parties are included in cabinet and so will give us more purchase on whether such mistaken beliefs are driven by PSI.[68]

Before we examine whether our estimated θ_{ij}'s predict partisan rationalization of this kind, it will be useful to first ask, as we did in the precious sections, if we can detect it using the traditional measure of PID. To examine this, we use our respondent-party dataset to estimate a logistic model in which the dependent variable indicates whether a respondent believed a given party was in the cabinet or not. Further, given that we strongly expect that more politically sophisticated individuals will be both less likely to be mistaken about which parties are in the cabinet and engage in less (of this kind of) partisan rationalization, we include an interaction between PID and the respondent's level of political sophistication (see Appendix D.6 for details on the construction of this sophistication index).

---

[66] We did not ask this question in the Netherlands because we fielded the survey during an on-going cabinet formation episode.
[67] The Italian and Danish cabinets formed, respectively, about 8 and 24 months before our survey.
[68] See Appendix F.34.

Table 16: Traditional PID and partisan rationalization: beliefs about which parties are in the cabinet

	(1)	(2)
Trad. PID	0.40*	0.43*
	(0.02)	(0.02)
R's Pol. Soph. Score		0.19*
		(0.02)
Trad. PID * Pol. Soph. Score		-0.17*
		(0.02)
Intercept	-1.12*	-1.13*
	(0.01)	(0.01)
Observations	32,932	32,932

The dependent variable is an indicator for whether the respondent thought the party was in the cabinet. The model is a logit. Numbers in parentheses are standard errors. All respondents from Denmark, Italy, Sweden, and the UK are included. Stars indicate the coefficient is significant at $p < .05$.

Table 16 shows that this particularly difficult form of partisan rationalization occurs with the traditional PID measure. Further, it interacts as expected with political sophistication: more politically sophisticated respondents evidence much lower levels of partisan rationalization.

Do we see this effect with our measure of PSI? Columns 1 and 2 of Table 17 report estimates of the same models as in Table 16 but replace traditional PID with our measure of PSI. Further, since we want to examine the hypothesis that positive partisans will be more likely to falsely report that the party they are positively attached to is in the cabinet, we only include cases of non-negative PSI in this table (we report the results for negative PSI in the next table). Finally, we also estimate the interactive model for traditional independents as well as a sample that includes only traditional partisans but excludes the party to which the respondent was (traditionally) attached.

Table 17: θ_{ij}^{+} and beliefs about which ***non***-cabinet parties are in the cabinet

	All R's (1)	All R's (2)	Trad. Inds. (3)	Trad. Partisans, PID = 0 (4)
θ_{ij}^{+} for non-cabinet party	0.34*	0.36*	0.26	0.38*
	(0.04)	(0.04)	(0.21)	(0.06)

	All R's (1)	All R's (2)	Trad. Inds. (3)	Trad. Partisans, PID = 0 (4)
R's Pol. Soph. Score		0.23*	0.39*	0.13*
		(0.04)	(0.11)	(0.05)
θ_{ij}^{+} * Pol. Soph. Score		-0.42*	-0.70*	-0.34*
		(0.05)	(0.19)	(0.07)
Intercept	-1.83*	-1.78*	-2.02*	-1.76*
	(0.03)	(0.03)	(0.14)	(0.04)
Observations	14,310	14,310	1,866	10,538

*DV is coded "1" if R indicated the party was in the cabinet and 0 otherwise. Data are respondent-party. Only parties that were **not** in the cabinet at the time of the survey are included. Model 4 includes only traditional partisans but excludes the party to which they are traditionally attached. The model is a logit. Numbers in parentheses are standard errors. Stars indicate the coefficient is significant at $p < .05$.*

The results in Table 17 are clear: In all models (including those run on samples including only respondent-party pairs in which traditional PID is zero), stronger positive PSI for a party that is not in the cabinet is associated with a higher probability that a respondent will falsely report that the party *is* in cabinet. Further, this association gets smaller as respondents are more politically sophisticated.

We also find large associations for traditional independents and traditional partisans (who are non-partisan with respect to all the parties included in the sample) that are substantively similar to those we get in our full sample. Thus, not only can our measure be used to draw the same conclusions about partisan rationalization as the traditional measure, but also to show that partisan rationalization, which is almost a litmus test for partisanship as a social identity, can be found among a much larger share of the electorate than previously thought (i.e., among traditionally measured non-partisans and independents).

Finally, Figure 22 uses the estimates in Table 17 to examine the substantive size of these effects and they are substantial.

[Figure 22 here]

Next, we can replicate the analysis above for negative partisans – though here it is less clear what we should expect. Should a person who identifies with the group of people who oppose a party (our definition of negative PSI) be less likely to think that party is in the cabinet, even when it is? If a negative PSI creates the same kind of partisan screen as positive partisanship, we might think so. That is not, however, what we find.

Table 18 and Figure 23 give the results for the same analysis as we did above for positive partisans, but instead with negative partisans. Specifically, the kind of "wishful thinking" we are looking for in this case is a positive relationship between a respondent's level of negative attachment to a cabinet party and the probability he or she incorrectly thinks the party is not in the cabinet.

Table 18: θ_{ij}^{-} and beliefs about which cabinet parties are ***not*** in the cabinet

	All R's (1)	All R's (2)	Trad. Inds. (3)	Trad. Partisans, PID = 0 (4)
$abs(\theta_{ij}^{-})$ for cabinet party	-0.67*	-0.37*	-0.51*	-0.38*
	(0.05)	(0.05)	(0.15)	(0.06)
R's Pol. Soph. Score		-0.80*	-0.82*	-0.78*
		(0.05)	(0.11)	(0.07)
$abs(\theta_{ij}^{-})$ * Pol. Soph. Score		0.23*	-0.03	0.33*
		(0.06)	(0.13)	(0.08)
Intercept	0.26*	-0.01	0.18	-0.06
	(0.04)	(0.05)	(0.14)	(0.05)
Observations	4,882	4,882	993	3,782

DV is coded "1" if R indicated the party was NOT in the cabinet and 0 otherwise. Data are respondent-party. Only parties that were in the cabinet at the time of the survey are included. Since we only use observations for which $\theta_{ij} \leq 0$, we take the absolute value of this variable, so that the signs of the coefficients can be interpreted the same way as those in Table 17 (so a positive coefficient in the first row would indicate this kind of partisan rationalization). Model 4 includes only traditional partisans but excludes the party to which they are traditionally attached. The model is a logit. Numbers in parentheses are standard errors. Stars indicate the coefficient is significant at p < .05.

These results are quite different from those in Table 17. While we see the same, expected, dependence of these beliefs on political sophistication (respondents are less likely to be wrong as they are more politically sophisticated), we get the opposite result than the one we found for positive attachment. That is, for all levels of political sophistication, being a stronger negative partisan of a party leads to *more* accurate beliefs about whether the party (to which one is negatively attached) is in the cabinet.

[Figure 23 here]

This result is intriguing and, along with some of the other results about negative partisanship discussed above, suggests that negative partisanship may be different in important ways from positive partisanship. Such differences in the empirical patterns for positive and negative partisanship are not uncommon in the literature, which is why Bankert (2024), in her study of negative partisanship, hesitates to fully embrace the position that negative partisanship is the same kind of social identity as positive partisanship. Given space limitations in this Element, we have not attempted a full exploration of negative partisanship in our data; however, results such as this one (contrasted with others in which negative partisanship does seem to work as the mirror of positive partisanship) suggest that the question of what negative partisanship really remains open.

6. Conclusions and the path forward

While more than fifty years of scholarship has demonstrated the myriads of ways that voters in multi-party systems behave differently than their counterparts in two-party systems, many students of comparative political behavior have continued to conceptualize partisan attachments in multi-party systems in exactly the same way as in two-party systems. In this Element, we argue it is time for that to change. Building on recent advances in the measurement of partisanship, we developed a multi-item battery of questions that can efficiently measure both positive and negative partisan attachment for a large number of parties. Using that measure, we find (for the first time) clear evidence that people hold multiple partisan attachments. Large majorities of respondents in each of our countries claim more than one party as "my party" and the sets of parties so identified are quite sensible (e.g., they are predictable from the parties' relative left-right positions, patterns of partisan cooperation, shared policy positions, and shared social group support). Further, these attachments have exactly the psychological consequences that we should expect if they really are phycological identifications of the kind usually studied using maximalist measures – e.g., observed levels of partisan rationalization across parties that are proportional to the strength of the voter's attachment to each.

Given this, we argue that multi-party attachments may help solve outstanding theoretical and empirical puzzles in comparative political behavior. These includes a potential explanation for the apparent (but often debated) decline in partisan attachments for some parties, countries, and time periods; as well as reconciliation of the oft-demonstrated importance of partisanship in two-party democracies with the long-standing critique that partisanship, in some multi-party democracies, is much less important – and is, indeed, less stable than the vote itself. As summarized in Chapter 1 and worked out in more detail in Appendix E, a very simple model of rational voting in multi-party contexts in which partisans are attached to multiple parties produces the hypothesis that the apparent instability of partisan identification relative to vote choice in places like the Netherlands may stem from the practice of measuring party identification with a maximalist, single-party question. That is, variability in a voter's reported *maximum* party identification from election to election may mask stability in her underlying vector of multi-party identifications. As we show in Appendix E, all that is required for this to happen is that the voter has at least two parties, *j* and *j'*, for which θ_{ij}^{+} and θ_{ij}^{+}, are close enough that short-term factors in the election (e.g., leader evaluations, economic evaluations, etc.) can easily overwhelm the difference between them.

While definitive tests of these ideas will require panel data on multi-party attachments (e.g., to test if they are indeed more stable than maximalist measures of PID), we hope this Element, by clarifying relevant theoretical concepts and demonstrating the efficacy of our measurement strategy, will increase the chances that these surveys will be done.

In the meantime, however, we do have some limited evidence from our cross-sectional data that provides empirical leverage on the question. Specifically, the historic levels of dissatisfaction with the Conservative party leading up to their defeat in the 2024 UK election creates an opportunity for us – because it was a particularly large shock in the kinds of "short-term factors" that impact both vote choices and (to a lesser extent – see Appendix E) partisan attachments. Specifically, in the theoretical model that we work out in Appendix E, we show that when partisans of a given party are

dissatisfied with that party's performance on relevant short-term factors and they have another party to which they are also attached, it is quite easy for relatively minor changes in a voter's underlying attachment vector to produce "switches" in both the vote and PID, even though the actual extent of change in the partisan attachment vector is small. This means that empirical analyses of the stability of PID relative to the vote, or of the impact of PID on the vote that use maximalist measures will necessarily reveal less stability (and impact on the vote) of PID than in two-party systems, even if the voters in both systems have the same levels of PID and, rely on it in vote choice, to exactly the same extent. Given the historically high levels of dissatisfaction among Conservative identifiers (including multi-party identifiers) with that party, the 2024 UK election is a particularly good election to expose such dynamics.

Specifically, one way to get a glimpse of whether underlying levels of multi-party partisanship were more stable than either the vote or maximalist measures of levels of PID, is to compare three ratios that we can calculate from our survey of the UK in 2024. First, we can calculate the ratio of Labour to Conservative identifiers using the traditional PID measure (call this "L/C traditional PID"). Next, we can calculate the ratio of the average strength of attachment to Labour vs the Conservatives for the same respondents using our measure ($\theta_{i\,LAB}^{+}/\theta_{i\,CON}^{+}$ or "L/C multi-party PID"). Finally, we can calculate the ratio of the actual vote for Labour vs. the Conservatives ("L/C actual Vote").

If, as our theory claims, short-term factors (which were strongly against the Conservatives) move reported maximalist measures of PID in the same direction as the vote (so that PID appears to "follow the vote") to a much greater extent than these factors move the actual underlying multi-party identities, then it should be the case that "L/C traditional PID" is much closer to "L/C actual vote" than is "L/C multi-party PID" – and that is exactly what we find.

The Labour to Conservative vote ratio (L/C actual vote) was1.42 compared to a Labour to Conservative PID ratio (L/C traditional PID) of 1.49. It is exactly this kind of correspondence – and the fact that the PID ratio tends to stay in sync with the vagaries of the vote from election to election – that fuels the narrative that PID moves with the vote in Europe (Clarke et al. 2009). If, however, we look at the ratio of $\theta_{i\,LAB}^{+}$ to $\theta_{i\,CON}^{+}$ (L/C multiparty PID) we find that it is just 1.07. That is, *our measure of underlying partisan attachment to Labour and Conservatives is considerably less favorable to Labour than the maximalist measure and much more separated from the reported vote*. This is consistent with the idea that movements in the underlying attachments of British Conservatives were amplified by maximalist measures of PID to produce greater apparent movement in partisanship than really happened.[69]

Besides highlighting the potential of multi-party identification to contribute to these ongoing theoretical debates in the literature, our other goals for this Element were to put the study of multi-party identification on a firm conceptual foundation, to work out the meaning and measurement of important concepts like total strength of identification and co-partisanship, to build a practical

[69] Without longitudinal data, we cannot say definitively that underlying multi-party attachments *moved* less than traditional PID towards the vote in this election; but this evidence is certainly consistent with such an interpretation.

empirical approach that can be implemented in surveys, and to explore its usefulness relative to traditional, single-party measures. In those efforts, we think we a have been successful, though there is certainly more to be done.

First, our multi-item, party attachment battery passes all the usual tests for a coherent, unidimensional scale. Further, it provides measures of both positive and negative PSI using only one scale, which makes measurement of both feasible in surveys with many parties. In addition, our survey respondents appear willing to engage with our question batteries in much the same way they do other matrix style questions common to election surveys (see Appendix A3.1.3). Thus, we think the extra survey time is a small price to pay for the rich data our measure provides about respondents' full partisan identities. In addition, as discussed in Appendix G, shorter question batteries perform nearly as well as our 5-question (balanced) question battery, as do unbalanced versions of the measure, and versions focused only on positive attachments.

That said, and despite the plethora of robustness checks and alternative measures of our underlying attachment vector that are provided in the appendices, we think of the measures of multiparty attachment provided in this Element as only a first step. There are important methodological questions that remain to be answered about how best to estimate multi-party attachments. Some of these include whether we give up too much in order to measure negative and positive attachment in the same battery; whether we should use compensatory multi-dimensional IRT models rather than the independent models we use here; and whether it is better to use party-specific models or to pool at the country level. Of course, such methodological questions are only worth answering if other scholars agree that the concept of multi-party attachments is needful and useful. We are convinced so far that it is, and we hope this Element is a first step in convincing others.

Third, we have shown that our estimated attachment vectors provide substantial predictive power over and above traditional measures of PID – focusing on the most probative set of respondents for judging this: traditional independents and traditional partisans who are non-partisans with respect to all but one of the parties. For vote choice, participation and engagement, partisan rationalization, and political efficacy, we find that our measures of attachment are strongly associated with outcomes among respondents that are otherwise indistinguishable using traditional minimalist measures of PID. Further, the predicted positive associations between the strength of partisanship and these attitudes and behaviors seem to hold across the full range of positive and (in many cases) negative identity strengths.

Fourth, we highlighted the important finding that the impact of partisan attachment on evaluations of a given party appears to depend mainly on the voter's affective attachment to ***that*** party, with cross-party impacts (i.e., my attachment to party A impacts my evaluations of party B) small and statistically insignificant.

Fifth, at least some of our results on negative partisanship (similar to Bankert's from 2024) suggest significant differences between positive and negative multi-party attachments and so challenge us to think carefully about the theoretical status of negative PSI and whether this is the right way to think about negative PID in general.

References

The list of references here are those we see as most central to our goals in this Element. See Online Appendix G for a full list of other work cited.

Achen, Christopher, and Larry Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press.

Adams, James, Simon Weschle, and Christopher Wlezien. 2021. "Elite Interactions and Voters' Perceptions of Parties' Policy Positions." *American Journal of Political Science* 65(1): 101–14.

Arian, Asher, and Michal Shamir. 1983. "The Primarily Political Functions of the Left-Right Continuum." *Comparative Politics* 15(2): 139–58.

Bankert, Alexa. 2024. *When Politics Becomes Personal: The Effect of Partisan Identity on Anti-Democratic Behavior*. Cambridge University Press.

Duch, Raymond, and Randolph Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge University Press.

Fortunato, David, Randolph T. Stevenson, and Greg Vonnahme. 2016. "Context and Political Knowledge: Explaining Cross-National Variation in Partisan Left-Right Knowledge." *The Journal of Politics* 78(4): 1211–28.

Garry, John. 2007. "Making 'Party Identification' More Versatile: Operationalising the Concept for the Multiparty Setting." *Electoral Studies* 26(2): 346–58.

Gidron, Noam, James Adams, and Will Horne. 2023. "Who Dislikes Whom? Affective Polarization between Pairs of Parties in Western Democracies." *British Journal of Political Science* 53(3): 997–1015.

Green, Donald, Bradley Palmquist, and Eric Schickler. 2002. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. New Haven and London: Yale University Press.

Holmberg, Sören, and Henrik Oscarsson. 2020. "Party Identification: Down but Not Out." In *Research Handbook on Political Partisanship*, eds. Henrik Oscarsson and Sören Holmberg.

Horne, Will, James Adams, and Noam Gidron. 2023. "The Way We Were: How Histories of Co-Governance Alleviate Partisan Hostility." *Comparative Political Studies* 56(3): 299–325.

Huddy, Leonie. 2001. "From Social to Political Identity: A Critical Examination of Social Identity Theory." *Political Psychology* 22(1): 127–56.

Huddy, Leonie, Lilliana Mason, and Lene Aarøe. 2015. "Expressive Partisanship: Campaign Involvement, Political Emotion, and Partisan Identity." *American Political Science Review* 109(1): 1–17.

Jost, John T. 2021. *Left and Right: The Psychological Significance of a Political Distinction*. Oxford University Press.

Kinder, Donald R., and Nathan P. Kalmoe. 2017. *Neither Liberal Nor Conservative: Ideological Innocence in the American Public*. University of Chicago Press.

Lee, Seonghui, Lie Philip Santoso, and Randolph Stevenson. 2025. *What Is the Left/Right in the Minds of Voters? An Experimental Analysis*. Cambridge University Press. Forthcoming.