

Mobilization, Persuasion, and the Partisan Fallout of the Gender Gap in U.S. Voting

Barry C. Burden* Michael G. DeCrescenzo†

October 31, 2016

This paper was prepared for presentation at 2016 Toronto Political Behavior Workshop. We thank Devin Judge-Lord and participants the American Politics Workshop at the University of Wisconsin-Madison for comments on an earlier draft.

*Professor of Political Science, University of Wisconsin–Madison (bcburden@wisc.edu)

†Graduate Student, University of Wisconsin–Madison (decrescenzo@wisc.edu)

Abstract

Since the gender gap in U.S. elections—in which women are more likely than men to vote for Democratic candidates—emerged in the early 1980s, observers have speculated that a larger gender gap works to the advantage of the Democratic Party. Any mechanisms linking the gender gap and the partisan vote share have yet to be elaborated let alone examined systematically. We argue that the size of the gap has no logically necessary connection to vote shares, and any connection that does exist depends on the mechanisms that produced the gap. We provide a theoretical framework in which an electoral gap is a function of two factors: mobilization and persuasion. We then elucidate the mechanisms behind the relationship using a two-stage model. In the aggregate, we find little evidence of a relationship between the size of the gender gap and the Democratic vote share in presidential elections since 1952. Additionally, when the mechanisms of mobilization and persuasion are traced over time, both mechanisms contradict the conventional wisdom. While the mechanism of mobilization has increased the gender gap, it has actually harmed the Democratic Party relative to the Republican Party. The mechanism of persuasion used to increase the gender gap, but the reduction in persuasion over time has helped the Democrats by increasing the partisan loyalty of Democratic voters. Finally, in affecting the Democratic vote through the gender gap, the joint impact of both mobilization and persuasion on the final vote share has decreased over time as the vote becomes a closer reflection of underlying party identification.

The percentage point gap in the vote choices of men and women has risen to double digits in recent U.S. national elections and has generated much analysis and discussion by the public, journalists, campaigns strategists, and academic researchers. Although the voting patterns of men and women are not as divergent as other subpopulations in the electorate such as blacks and whites, the gender gap is uniquely important because the two groups in question are nearly equal in size. Even small shifts in the vote choices of men or women could have acute effects on election outcomes.

We ask two research questions about the gender gap in voting in U.S. elections, the first more easily answered than the second. First, does the gender gap advantage the Democratic Party? Despite being widely accepted as folk wisdom, we are not aware of any empirical testing to show whether a bigger gap in the vote choices of men and women actually helps Democratic candidates and hurts Republicans. Second, to the extent that the gap affects the vote, what are the mechanisms that make it happen and how do they vary over time? In particular, we explore how much of the gap's impact is due to party supporters being *mobilized* versus the degree to which voters are *persuaded* to support one party over the other.

In response to the first question, we find little evidence of a relationship cross-sectionally or longitudinally between the size of the gender gap (defined as the percent of women who vote Democratic minus the percent of men who vote Democratic) and the overall Democratic vote share. Rather than assume a timeless "effect" of the gender gap on the vote, we argue that the consequences of the gender gap depend on precisely how the gap was generated. Without an understanding of the mechanisms that produce it, the gender gap's bearing on the final vote is impossible to extract.

In response to the second question, we develop a theory to describe how the gender gap can advantage one party over the other. We conceive of election outcomes as a function of mobilization (i.e., turnout of loyal partisans) and persuasion (i.e., defection of partisans to vote for the opposite candidate)—processes that may work separately among men and women and that vary over time. Using American National Election Study data over 16 presidential elections,

we apply a two-stage model to show how mobilization and persuasion among men and women affect both the gender gap and the Democratic vote.

Although the conventional wisdom assumes that a larger gender gap benefits the Democrats, we find that the effects of both mobilization and persuasion contradict the common narrative. Over time, the Democratic Party's advantage in mobilizing partisan voters has waned, mostly due to the growing Republican Party base especially among men. The growing mobilization of Republican men relative to Democratic men has increased the gender gap even as it has weakened the Democratic vote. Persuasion provided a strong benefit to Republican candidates in the past, especially in convincing Democratic men to cross the aisle and cast Republican votes. As swing-voting has declined in recent decades, however, so has its impact on the gender gap and its net benefit to Republicans. As a result, Democrats have benefited from declining persuasion even though declining persuasion has reduced the gender gap, all else equal. Lastly, although Democrats historically have benefited from mobilization and Republicans from persuasion, these advantages have declined in the era of polarization. We thus reach the ironic conclusion that just as awareness and interest in the gender gap have risen since the 1980s, polarization and sorting of the electorate have rendered it less impactful in determining the outcomes of elections.

1 THE GENDER GAP AND THE DEMOCRATIC VOTE

Observers generally agree that widespread interest in the gender gap in voting took hold after the 1980 presidential election when it became clear from exit polls that men chose Republican Ronald Reagan at a rate about eight points higher than that of women, a significant change from the previous election in which men and women voted almost identically (Norrander 1999; Wirls 1986). The persistence of this new cleavage between men and women spawned an array of academic studies seeking to understand the factors behind gender differences in party politics. From this work we have learned a great deal about how policy preferences, economic

circumstances, social position, and other experiences contribute to gender differences in political orientation (Box-Steffensmeier and Lin 2004; Chaney, Alvarez, and Nagler 1998; Conover 1988; Edlund and Pande 2002; Kaufmann and Petrocik 1999; Manza and Brooks 1998).

As rich as this literature has become, it tends to blur what might be an important distinction between party identification and voting behavior. Although gender gaps appear in both, research demonstrates that the two are not synonymous. Party ID is an attitude, which may be held weakly or intensely, whereas vote choice is a discrete decision that can either reflect party loyalty or run against it. The former also moves slowly over the long-term while the latter can fluctuate more rapidly in the short-term. Indeed, two general reasons why election results do not perfectly reflect the distribution of party identification in the electorate are that rates of turnout vary between the parties and because partisans cross party lines at different rates.¹ We pay close attention to this distinction as a means to understand the degree to which the gender gap in election returns is a simple function of underlying partisan preferences versus short-term campaign-specific factors that affect rates of turnout and defection differentially for men and women.

Journalists and pundits have widely assumed that Democratic candidates benefit when the vote choices of men and women diverge. Press coverage leading up to each election typically argues that Democrats are more successful at the ballot box when their advantage in the women's vote is greater. For example, in the run up to the 2012 presidential election, the *Washington Post* reported that "The Republicans are on the defensive partly because the gender gap—in which Democrats have a sizable advantage among women—is growing."² Turning the discussion around, observers often link Republican electoral successes to decreases in the size of the gender gap. A discussion of the parties' demographic bases in *The New York Times*, for example, concluded that Republicans succeeded in the 2014 elections in part because they refuted the Democrats' accusation of a "war on women" that "generated a gender gap and has been cru-

¹This insight is the basis of the Michigan school's classification of presidential elections as "maintaining," "deviating," or "realigning" (Campbell 1966) and the development of the concept of the "normal vote" based on party ID (Converse 1966).

²Robert McCartney, "Romney's Woes Threaten McDonnell and Allen," *Washington Post*. October 7, 2012.

cial to past Democratic victories.”³ Another report in *USA Today* on the eve of the Republicans’ triumph in the 2010 midterm elections observed that “men generally tilt toward GOP candidates, so a significant narrowing of the gender gap among women is contributing to Democrats’ struggles,” with one commentator concluding that Democrats would only win “by mobilizing women voters.”⁴ The discussion also frequently assumes that female voters are more volatile or persuadable, as suggested by popular images of “soccer moms” and “security moms” who are thought to be more susceptible to campaign effects than men. Only rarely does a media voice suggest that the political preferences of women and men do not differ enough for the gender gap to matter and that “You always know the Democrats are in big trouble when the media starts harping on the gender gap.”⁵ Aside from these rare exceptions, the folk wisdom remains that Democrats benefit from a wider gender gap.

Despite the prevalence of the conventional wisdom, its proponents have not offered a clear explanation for the relationship or provided even basic empirical evidence that it exists. In particular, the reasons are muddled as to why the gender gap and the vote might be related. Journalistic coverage has sometimes suggested that the key mechanism is the distribution of party loyalties among men and women, while other coverage has suggested that differential turnout by men and women is the main driver. We argue that both mechanisms are potentially important but that neither implies a specific or constant relationship between the gender gap and vote share.

We start by considering underlying partisanship in the electorate as a culprit. Figure 1 displays the gender gap in voting from 1952 to 2012 from the ANES. Here we examine the share of men and women who reported voting for one of the two major party candidates. Residue of the underlying gaps in party identification is apparent. Men are slightly more likely to vote Democratic in the 1950s, after which the differential collapses and then reverses to become the familiar contemporary gap in which the women’s vote is decidedly more Democratic. A visible

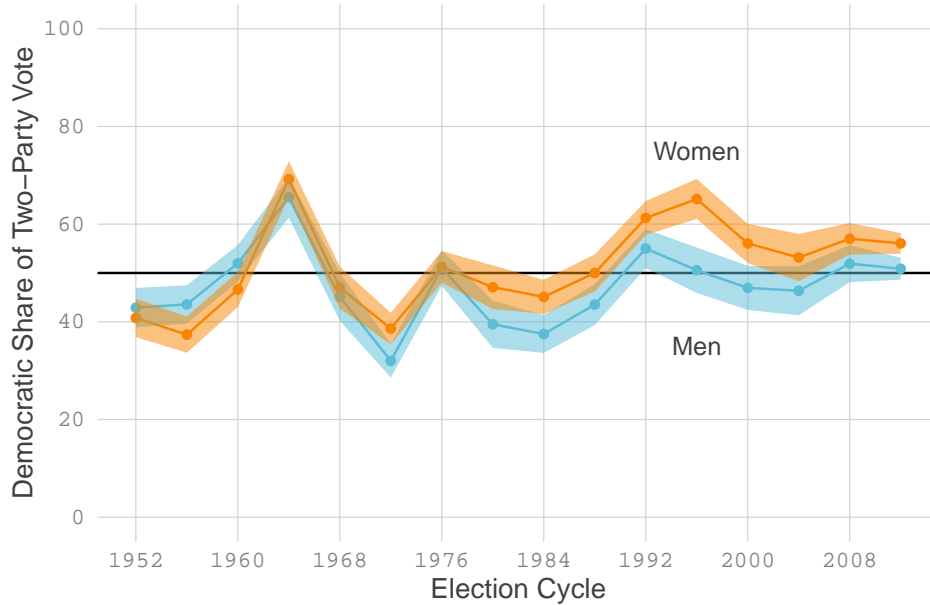
³Thomas B. Edsall, “Making the President Small,” *The New York Times*, November 5, 2014.

⁴Mimi Hall, “Democrats Lose Share of Women to GOP,” *USA Today*, October 20, 2010.

⁵Linda Hirschman, “Forget the Ladies; The Gender Gap Doesn’t Matter—and It Never Has,” *Newsweek*, November 5, 2012.

gap emerges briefly in 1972 only to reestablish itself in 1980 and then reach double digits in the 1990s. This shift among men in the 1980s is a first hint that the relationship between the gender gap and the partisan vote outcome is likely to vary over time rather than follow a universal pattern.

Figure 1 Democratic Voting of Men and Women in Presidential Elections




Source: ANES Cumulative File

Figure 2 shows the percentages of men women identifying with the Democratic and Republican parties in American National Election Study (ANES) in presidential election years from 1952 through 2012. Partisanship among women since the 1950s has been relatively stable, with between roughly 50 and 60 percent of women identifying with the Democratic party and roughly 30 to 40 percent of women identifying as Republicans. Among men, however, there has been a pronounced decline in Democratic Party identification—from 57 percent in 1952 to 43 percent in 2012—and a corresponding increase in Republican Party identification—from 32 percent in 1952 to 42 percent in 2012. This reflects the view of Kaufmann and Petrocik (1999) who argue that the gender gap in voting is caused not by a long-term strengthening of Democratic loyalty among women but by a long-term *weakening* of Democratic partisanship among

men. However, it appears that the pattern changed in the 1980s when the movement of women toward the Democrats was more consequential (Kaufmann 2006).

Figure 2 Party Identification of Men and Women




Source: ANES Cumulative File

Although party identification provides important raw material to form the gender gap in voting, the two are not the same. Kaufmann and Petrocik show that party ID often accounts for less than half of the voting gap. Following their insight, we argue that the observed gap in voting behavior must be a function of party ID in the larger electorate, filtered in two ways. The two parts of this function are how many partisans of each sex turn out to vote for their preferred parties (the *mobilization* effect) and how many defect to vote for the opposing party (the *persuasion* effect).

Differential turnout of men and women seems a likely candidate to explain the connection between the gender gap and the Democratic vote. The rise of the gender gap in both partisanship and voting coincides with differential trends in voter turnout among men and women. Figure 3 shows self-reported voter turnout among men and women in the ANES. Though turnout among women lagged behind men by roughly 10 points in the 1950s and early 1960s, the gap

in turnout has shrunk and even reversed itself in recent years, with women’s turnout anywhere from two to six percentage points higher than men’s turnout since 2004.⁶ As a result, rising turnout among women relative to men may produce more Democratic *votes* on net, even when holding the vote choices of men and women constant. Indeed, even though party identification among men stabilized around the 1988 presidential election, [Zingher \(2014\)](#) finds that women have nonetheless grown as a share of the Democratic Party’s electoral coalition since 1988.⁷

Figure 3 Self-Reported Voter Turnout in Presidential Elections



Source: ANES Cumulative File

An election with a larger gender gap might reflect a context favoring the Democratic candidate in which women—who are more likely than men to identify as Democrats—are more enthusiastic than men and end up voting at higher rates. This seems to be the logic behind Democratic campaign efforts to “get out the women’s vote.” The cross-pressures felt by men—who tend more Republican—could lead to lower levels of voter turnout for them. In contrast, in a setting where the Republican candidate is succeeding, Democratic-leaning women might

⁶Self-reported voter turnout data from the U.S. Census Bureau’s Current Population Survey (CPS) also show that women’s turnout has surpassed men’s, though the CPS observes this reversal as early as the 1980s.

⁷We note that during this time felon disenfranchisement laws have made an increasing share of the electorate ineligible to vote and that the voters affected are overwhelmingly male ([Manza and Uggen 2006](#)). This could be a partial explanation for why female turnout surpassed male turnout.

vote at lower rates because of their ambivalence while Republican-leaning men would participate more enthusiastically. This simple response, based on mobilizing effects of partisan attachments, could produce a positive correlation between the gender gap and the Democratic vote share. Whereas equal turnout among men and women would merely reflect the underlying gap in party identification, a turnout differential between men and women in each party would exaggerate or diminish the gender gap. This scenario requires some stringent assumptions, particularly the notion that each additional woman mobilized to vote is at least as Democratic as women already intending to vote, an unlikely scenario given research showing that regular voters have stronger partisan attachments than less regular voters (e.g., [DeNardo 1980](#)). Yet if these conditions hold, it would be wise for Democratic campaigns to focus on increasing turnout among women and for Republican campaigns to focus on increasing turnout among men.

Figure 3 hints that turnout affects the gender gap, but in a surprising way. As documented earlier, the gender gap has increased over this time period, rising from a negative gap in the 1950s to a more stable difference of five to 10 points by the 1990s. During this same time voter turnout among women increased not much at all (a correlation with time of $r = .06$, driven mostly by the last two elections in the series). In contrast, men's turnout fell over time ($r = -.65$). These patterns do not conform with the popular view that rising turnout among women is responsible for growth in the gender gap.


2 DOES THE GENDER GAP HELP THE DEMOCRATS?

Later in the paper we develop of a theory that incorporates both differential mobilization and differential persuasion of men and women. Our theory makes clear that the gender gap may either help or harm Democrats. Before detailing that logic, we first test for a global relationship between the gender gap and the vote. Although views about the gender gap are not monolithic, it is generally assumed that a wider gap corresponds to a larger share of the vote going

to Democrats. In this section we examine this hypothesis cross-sectionally and longitudinally. The test is simply whether there is a positive correlation between the size of the gender gap and Democratic vote.

For the cross-sectional test, we rely on exit polls from the 2004 and 2008 elections.⁸ Exit polls are conducted by a media consortium of five major television networks and the Associated Press. Because the two elections we examine are so different—2004 being a modest Republican victory and 2008 being a substantial Democratic victory—any conclusions that apply to both would seem generalizable to other recent election years. For each state we compute the gender gap as the percentage of women who voted for the Democratic presidential candidate minus the percentage of men who did so. In most cases this variable should be positive. The key outcome of interest is the official percentage of the vote won by the Democratic ticket. Figure 4 displays the scatterplots for the two election years. We have imposed the OLS regression line with 95% confidence intervals. The figure shows little sign of a consistent relationship, as evidenced by the positive but insignificant correlations of .16 in 2004 and .14 in 2008.

Figure 4 State-level Gender Gap and Democratic Vote

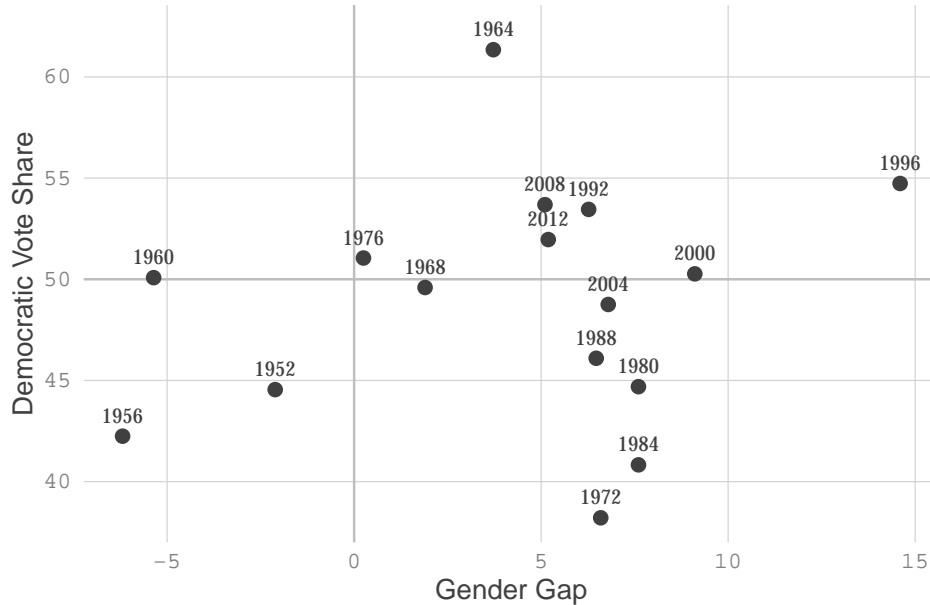


A second way to evaluate the relationship is over time. This approach has the benefit of cov-

⁸Adding data from 2012 could be useful, but exit polls were not conducted in 19 states that year.

ering a wider range of electoral contexts from landslide elections such as 1972 to close contests such as 2000. However, a limitation of the approach is that the ANES time series only covers 16 presidential elections, so our statistical power is limited. We again compute the Democratic share of the two-party vote as the outcome of interest and compute the gender gap as the difference between the percentages of women and men voting Democratic. Figure 5 displays the scatterplot for elections from 1952 to 2012. The data tell essentially the same story as the cross-sectional analysis. The correlation between the gender gap and the Democratic vote is a positive but insignificant .17 and the regression coefficient is a mere .18.

Figure 5 Gender Gap and the Democratic Presidential Vote



Based on these two sets of analyses testing the first research question, we conclude that an expansion of the gender gap has little direct bearing on the Democratic vote. Upon further examination, it will become clear why this is the case. The gender gap as traditionally defined is extremely useful for summarizing how the voting preferences of men and women compare in a particular election, but on its own the size of the gap does imply anything in particular for the shares of the vote won by the major party candidates.

In the remainder of the paper we explore how voter mobilization and persuasion—the two

primary effects of campaigns—contribute to the gender gap and mediate between the gap and the vote. We argue that an understanding of how the gap increases the Democratic vote share must account for changes not only in the party affiliations of men and women over time, but also the level of voter turnout and cross-party voting among men and women. In the following section, we provide a framework for understanding the impact of these electoral forces—particularly *mobilization* and *persuasion*—in both the formation of a gender gap in voting and the gap’s implications for the final vote outcome.

3 A THEORY OF ELECTORAL GAPS

Our analytical task becomes much more difficult when we move beyond simple “gaps” that appear on the surface between men and women—or any other groups in the electorate—to a deeper understanding of the mechanisms that produce them. The gender gap is conventionally defined as the Democratic vote share among women minus the Democratic vote share among men.⁹ This is how we measured the gap in the previous section of the paper. However, because vote shares are proportions, and because vote shares can be placed on the same scale regardless of *how many votes* they are worth, the gender gap by itself does not convey enough information to determine whether it is “good” or “bad” for Democrats. For example, if the Democratic vote share among women is 55 percent and the Democratic vote share among men is 45 percent, the resulting 10 percentage-point gap does not itself convey whether Democrats have any more *votes* than the Republicans do. We thus turn our focus to the mechanisms that cause voting gaps to form between groups in the electorate in the first place. Although our application is to the gender gap—a demographic cleavage that we believe is especially consequential because the rough parity and sheer size of the two voting blocs—the theoretical framework we offer can in principle be applied to any electoral gaps.

We contend that three primary factors shape voting gaps between groups in the electorate.

⁹There are alternatives such as [Norris \(2003\)](#) measure, which computes the net partisanship of women minus the net partisanship of men.

These factors are (1) the long-term *partisan inclinations* of group members, (2) the immediate *mobilization* of those party supporters in the group to vote, and (3) the short-term *persuasion* of group members to vote against their long-term partisan inclinations.¹⁰ While past scholarship has examined the role of partisanship in the formation of the gender gap (Carroll 2013; Kaufmann and Petrocik 1999), studies of the gender gap have not included mobilization and persuasion in a unified framework.

Our approach is reminiscent of Axelrod (1972)’s analysis of coalition composition, which measures a group’s contribution to a party coalition as the group’s *size* × *partisan loyalty* × *turnout* (see also Stanley and Niemi 1986). We differ in considering vote choice not as simple “loyalty” but as a combination of long-term partisan inclinations realized through mobilization and short-term partisan defections. We believe that our approach mimics how campaign strategists think about building winning coalitions by turning out supporters who have long-term attachments to the party (GOTV) and persuading supporters of the opposing party to defect in the short term.

We present a series of illustrations to outline our approach, building upward from micro-foundations to aggregate results. As shown in Table 1, if we assume a two-party system, a citizen must make one of three potential choices in the election: voting Democratic (D), voting Republican (R), or not voting.¹¹ This trichotomous choice exists for all groups in the electorate.

Table 1: Three Choices for Potential Voters

Do Not Vote	Vote D
	Vote R


¹⁰Group *size* is a factor that determines how group differences in voting affect the final vote outcome. Smaller groups have less influence on the final vote, but group size does not itself influence the size of a voting gap across groups. If Democrats win 50 percent of votes from Group A and 30 percent from Group B, there is a 20-point gap regardless of which group is numerically larger. Because the numbers of adult men and women in the U.S. are essentially equal, we do not explicitly incorporate group size into our analysis. Studies of other unequally-sized groups in an electorate (such as racial or ethnic groups) would have to take into account the relative sizes of the groups when determining the impact of group differences in voting on the final vote.

¹¹As the table suggests, not voting is typically the modal choice in actual elections. Voters and candidates who are not affiliated with the major parties are omitted.

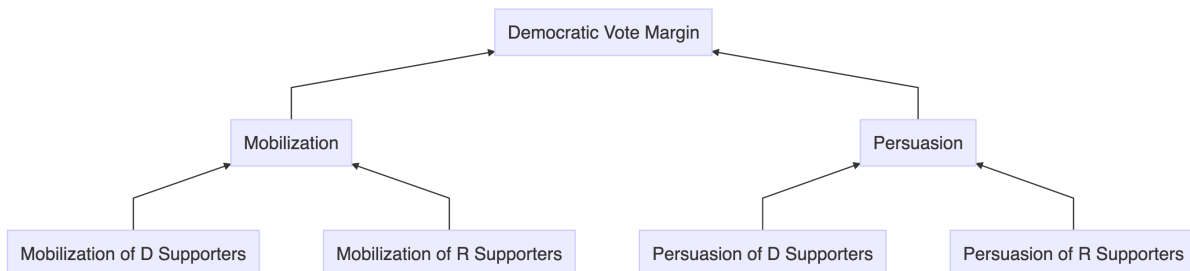
Now consider how these choices aggregate to form the gender gap. We assume there are only two groups in the electorate. The partisan voting gap between Group W (Women) and Group M (Men) is expressed as the difference in the percentage of women and men who vote D , provided that they do vote.¹²

$$\text{Gender Gap} = \% \text{ Vote } D_W - \% \text{ Vote } D_M \quad (1)$$

The level of party identification within each group is a serviceable baseline for the way these groups will vote—most partisans, if they vote, cast votes for their own party.¹³ But to obtain the final vote, we must account for the processes that intervene between partisanship and the vote: *mobilization* and *persuasion*. When we take these factors into account, each of the *Vote D* terms in Equation 1 is the result of both the mobilization of Democratic partisans and the persuasion of Republican partisans to crossover and vote Democratic.



The final Democratic vote margin is therefore the result of mobilization and persuasion mechanisms among Democrats, net of those same mechanisms among Republicans.



¹²This is a percentage of the “two-party vote.” Because the gender gap in partisanship is partly due to different rates of independent affiliation for men and women (Norrander 1997), future work might incorporate voting for non-major party candidates.

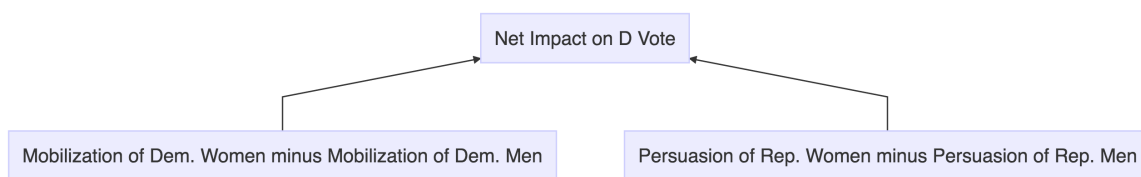
¹³Other measures of party support or affiliation are possible. We take party ID as a fair approximation of a person’s default or “standing position” in the absence of other forces.

Two points are worth noting. First, persuasion has twice the impact on the vote margin that mobilization does. This is because mobilizing one additional Democratic-aligned voter nets the Democratic Party *one* vote, whereas persuading a Republican-aligned voter to vote Democratic nets the Democratic Party *two* votes—Democrats gain a vote and Republicans lose one. Although persuasion might be more difficult for campaigns than is GOTV, it has twice the payoff. Second, analysts must be cautious when measuring mobilization and persuasion as *rates*—such as turnout rates or rates of cross-party voting. If they do not account for differing denominators, simple rates do not capture the real numerical impacts of mobilization and persuasion on the final vote. Democratic-aligned citizens and Republican-aligned citizens may have identical turnout rates, for example, but if a greater number of citizens is aligned with the Democrats than with the Republicans, Democrats receive a greater number of votes, all else equal. Because the impact of the mobilization of supporters depends on both the number of supporters and their rate of turnout, it is helpful to conceive of the *mobilization* component as distinguishing turnout among partisans from underlying partisanship in the electorate. The number of mobilized partisans is the total number of partisan supporters minus the nonvoting partisans.¹⁴ Equation 2 shows this logic.

$$\text{Mobilization of D supporters} = (\text{Total D supporters}) - (\text{Nonvoting D supporters}) \quad (2)$$

By conceiving of mobilization and persuasion in a given election as a raw total rather than as a rate, we can account for mobilization and persuasion among men and women separately to measure more directly how gender differences in voting behavior affect the partisan outcome of elections. The impact of women’s voting behavior on the vote, net of men’s voting behavior, can be understood as the following:

¹⁴The *mobilization* concept (e.g. *Mobilization of D supporters*) will overestimate the number of votes for a party, since not all party supporters will vote loyally. However, this over-estimation is offset by a corresponding persuasion term (e.g. *Persuasion of D supporters to vote R*) in a full equation.



This formulation makes clear that relationship between the gender gap and the Democratic vote could be the result of differential turnout among party supporters, differential persuasion of those supporters, or both. Though past studies have examined the role of voter turnout and cross-party voting in the determination of election outcomes (Hillygus 2014; Huber and Arce-neaux 2007), our approach is unique in that we decompose the entire vote into the sum of these two mechanisms.

4 MEASURING NET MOBILIZATION AND PERSUASION

To model the Democratic vote as the sum of persuasion and mobilization, we need measures of each concept that will accurately capture each mechanism’s numerical impact on the vote across time. Conventional measures of turnout and defection rates for each gender in each party are inadequate for this purpose because those rates are measured over different denominators. For example, both Democratic and Republican men may have a turnout rate of 60%, but if the pool of Republican men is greater, turnout rates will not capture the impact of partisan mobilization on the vote. The rate of cross-party voting is vulnerable to the same measurement problem. Measures of mobilization and persuasion, therefore, require a strategy that can capture the numerical impact of these mechanisms on the vote across gender, across party, and across time. In this section, we explain how our measurement approach satisfies these criteria.

We create four measures to capture the partisan impact of mobilization and persuasion separately for men and women. Each of these measures is designed to capture the *net partisan impact* of each mechanism for each gender. In other words, these variables measure *which*

party is the net beneficiary of mobilization among women, mobilization among men, persuasion among women, and persuasion among men—and by how much.

To capture the net partisan impact of mobilization, we create measures of the *net Democratic advantage in turnout* among men and women. We measure mobilization for a party as the number of partisans who reported voting in an election divided by total number of eligible voters in that year. We compute this quantity separate for men and women. For example, we measure the mobilization of Democratic men as:

$$\text{Mobilization of Democratic Men} = \frac{\text{\# of Male Democratic Voters}}{(\text{Total \# of Voters and Nonvoters})}$$

The result is a measure that represents the *share of the total electorate* that is composed of Democratic men who turn out to vote. We repeat this procedure for each gender-by-party combination for every presidential election year.¹⁵ This measurement strategy captures the numerical impact of the mobilization on the final vote because the denominator for each election cycle is equivalent across men and women and across Democrats and Republicans, and we can compare the impact of mobilization across election years because we divide by the total number of voting-eligible adults in the ANES for each election cycle.

We then calculate the *net Democratic advantage in turnout* for each gender by taking the mobilization measures for Democrats minus the mobilization measures for Republicans. For example, the net Democratic turnout advantage among men is:

$$\text{Mobilization of Democratic Men} - \text{Mobilization of Republican Men}$$

The measure captures the *net Democratic advantage* because the quantity is zero when an equal number of Democratic and Republican men turn out to vote,¹⁶ and it is positive when

¹⁵This formulation excludes independents by necessity, as we can make no uniform judgment as to whether turnout among independents should advantage either party *a priori*.

¹⁶That is, when the *share of the eligible electorate* is equally composed of Democratic and Republican men who turned out.

Democrats have the numerical advantage. We construct a similar measure for women’s mobilization.¹⁷

We follow a similar strategy to operationalize the *net Democratic advantage in persuasion*. Because the *percentage* of defecting partisans will have a greater or lesser impact on the final vote depending on the size of the partisan base, we implement the same correction as above by dividing the number of party defectors by the total number of voting-eligible adults. This keeps the mobilization and persuasion measures on the same scale to compare their relative impacts on the gender gap and Democratic vote share. Persuasion among Democratic men, for example, is thus measured as:

$$\text{Persuasion of Dem. Men} = \frac{\# \text{ of Dem. Men Voting Rep.}}{(\text{Total \# of Voters and Nonvoters})}$$

We again compute the *net Democratic advantage in persuasion* by taking a difference across parties. The *net Democratic advantage in persuasion* among women, for example, is therefore:

$$\text{Persuasion of Rep. Women} - \text{Persuasion of Dem. Women.}$$

As with the net mobilization measures, the net persuasion measures are positive when Democrats are the numerical beneficiaries of persuasion.

Models that incorporate these measures are estimated two different ways to reflect alternative operationalizations of party ID: one based only on “strong” and “weak” partisans and another that also codes independent “leaners” as partisans. The scholarly literature is divided as to whether “leaners” ought to be considered independent or are better viewed as partisans because of their voting behavior (Keith et al. 1992; Miller and Shanks 1996; Petrocik 2009). Examining both conceptualizations of partisanship is important theoretically because it allows us

¹⁷By taking differences to calculate the *net Democratic advantage in mobilization* among men and women, we are able to cut the number of variables in half (from four to two) and save degrees of freedom in subsequent regressions while losing no information about the numerical advantage either party has over the other in turnout in a given election year.


to separate mobilization more clearly from persuasion by addressing the concern that “leaners” are independents who report their vote intentions when they admit that they are “closer” to one party than the other. We present the results for both codings. Fortunately our results do not depend on which measure is used.

Figure 6 summarizes these descriptive measures of net mobilization and persuasion over time, with independent leaners coded as independents (in the upper panels) and as partisans (in the lower panels). The mobilization panels remind us that in the earliest years of the ANES, soon after New Deal era and World War II, larger shares of the eligible electorate were composed of Democratic-affiliated men and women than Republican-affiliated men and women. These advantages were offset to a considerable degree by Republican advantages in *persuasion* among men and women—it is important to keep in mind that persuasion is worth double the votes of mobilization in the final election result. This means that in many of these early survey years, although a net advantage in mobilization for Democrats appears to be substantial, this mobilization is outweighed by a persuasion advantage for Republicans, hence Republican victories in these elections despite the overwhelming advantage in partisanship for Democrats (Shafer and Claggett 1995).

When shifting to more recent survey years, we note two important trends in net mobilization and net persuasion that implicate both the gender gap and the Democratic vote. First, the Democratic Party’s advantage in the number of mobilized men shrinks over time and even reverses occasionally in recent contests, while the Democratic advantage in mobilization among women remains positive and sizable throughout the time period under consideration. This reflects the long-term partisan drift among men toward the Republican Party and the relatively stable Democratic majority in party identification among women, as originally highlighted by Kaufmann and Petrocik (1999).

Second, although early survey years provided Republicans with a sizable benefit from persuading Democrats to cast Republican votes, this persuasion advantage shrinks substantially beginning in the late 1980s and early 1990s. By 2008 and 2012, neither party receives a sizable

Figure 6 Net Democratic Advantage in Mobilization and Turnout, 1952–2012



advantage in persuasion. This is reminiscent of other findings from the literature on voting behavior in the contemporary era of partisan polarization, with high levels of party-loyal voting and relatively little swing voting in recent election cycles (Bafumi and Shapiro 2009; Bartels 2000).

5 CONNECTING THE MECHANISMS IN A TWO-STAGE MODEL

Using the net mobilization and net persuasion variables described in the previous section, we can now study how these mechanisms interact over time to produce the familiar gender gap in voting. Additionally, we can uncover how the formation of the gender gap in each presidential election cycle works to the benefit or the detriment of the Democratic Party. To accomplish this, we employ a regression approach that proceeds in two stages. The technique is a bit unusual,

but it similar to the more conventional two-stage least squares (TSLS) procedure in which the predicted outcome from a first-stage model is included as an explanatory variable in a model of the second-stage outcome. In our case, the TSLS approach would predict the gender gap in stage one using net mobilization and persuasion among men and women as predictors, and it would proceed to use the estimated gender gap to predict the Democratic vote in the second stage. Because the second-two equation would include the estimated gender gap as a single predictor, however, information about how each individual mechanism affects the Democratic vote *through* the predicted gender gap would be lost. To avoid losing information about the unique impact of each predictor, we instead decompose the predicted gender gap from stage one to reflect the contributions of the four mechanisms.

To illustrate, if the gender gap in an election year (G_t) is estimated with net mobilization and persuasion among men (subscripted m) and women (subscripted w), the predicted gender gap is

$$\hat{G}_t = \hat{\alpha} + \hat{\beta}_1 (\text{Mob}_{m,t}) + \hat{\beta}_2 (\text{Mob}_{w,t}) + \hat{\beta}_3 (\text{Per}_{m,t}) + \hat{\beta}_4 (\text{Per}_{w,t}). \quad (3)$$

Whereas a typical second-stage regression would include the gender gap estimated from stage one as a predictor of the vote (V_t), we decompose the estimated gender gap into its constituent mechanisms with the following procedure.

Each mechanism's partial impact on the gender gap is summarized by its estimated coefficient from stage one. For example, if $\hat{\beta}_1$ captures the partial effect of mobilization among men on the gender gap, then $\hat{\beta}_1 (\text{Mob}_{m,t})$ captures the effect of men's mobilization on the gender gap in every year t (with other factors held constant). We set up a second-stage equation to estimate how each mechanism's *impact on the gender gap* is related to the final Democratic vote. The second stage includes the $\hat{\beta}_1 (\text{Mob}_{m,t})$ term as a predictor in a regression that predicts the Democratic vote share, and the estimated coefficient for this term in stage-two would show how the impact of men's mobilization on the gender gap translates into an impact on the final vote.

We include each mechanism and its estimated coefficient from stage one as predictors in the stage-two regression to predict the officially reported Democratic vote share (V_t):

$$V_t = \hat{\gamma} + \hat{\delta}_1 (\hat{\beta}_1 \text{Mob}_{m,t}) + \hat{\delta}_2 (\hat{\beta}_2 \text{Mob}_{w,t}) + \hat{\delta}_3 (\hat{\beta}_3 \text{Per}_{m,t}) + \hat{\delta}_4 (\hat{\beta}_4 \text{Per}_{w,t}), \quad (4)$$

where $\hat{\gamma}$ is a constant, each $\hat{\beta}$ term is the estimated coefficient from the stage-one regression, and each $\hat{\delta}$ term measures how each mechanism's effect on the gender gap is related to the final vote share. Because the underlying data varies with t , we avoid making the assumption that there is a timeless relationship between the gender gap and the vote. Instead, the extent to which the gender gap is related to the Democratic vote is a function of the extent to which each underlying mechanism (mobilization and persuasion) nets the Democratic Party a greater or lesser number of votes than the Republican Party.

Table 2 outlines our expectations for the two-stage procedure. The first stage (left) predicts the gender gap with the net mobilization and persuasion predictors. Because the gender gap measures the percentage difference in the Democratic vote between men and women, all else equal, the gender gap should *increase* when the Democratic vote among women increases relative to the Democratic vote among men, just as it should increase when the Republican vote among men grows relative to the Republican vote among women. Conversely, the gap should *decrease* when the Democratic vote among men grows and when the Republican vote among women grows, all else held equal. Thus, we would expect a net Democratic advantage in mobilization and in persuasion among women to have a positive impact on the gender gap, while we expect a Democratic advantage in these mechanisms among men to exert a negative impact. And as explained above, the persuasion measures should net the advantaged party two votes rather than one; we thus expect the magnitude of persuasion coefficients to be approximately double the magnitude of mobilization coefficients.¹⁸

The columns for stage two in Table 2 carry these mechanisms through the gender gap to

¹⁸This two-to-one relationship will not be perfect in part because we are excluding independents from the persuasion measure.

Table 2: Hypothesized effects of net mobilization and persuasion in two-stage model

Mechanism	Stage one (Gender Gap)		Stage two (Democratic Vote)	
	Direction	Magnitude	Direction	Magnitude
Net Dem. Mobilization – Women	+	1	+	1
Net Dem. Mobilization – Men	–	1	–	1
Net Dem. Persuasion – Women	+	2	+	2
Net Dem. Persuasion – Men	–	2	–	2

the final vote. Net mobilization and persuasion among women increase the gender gap by increasing the number of Democratic votes garnered from women, so to the extent that these mechanisms increase the gender gap, they also increase the Democratic vote share. We therefore expect the coefficients for the net Democratic mobilization and persuasion of women to be positive. On the other hand, we know that Democratic mobilization and persuasion among men should reduce the gender gap. As such, when net mobilization and persuasion among men advantage the Democratic Party (and thus take positive values), this means that a smaller gender gap produces a larger Democratic vote share. So to the extent that the mobilization and predictors exert an *upward* force on the gender gap in stage one, this indicates that Republicans gain more votes from men than Democrats. We thus expect the stage-two predictors for mobilization and persuasion among men to be negative as well.

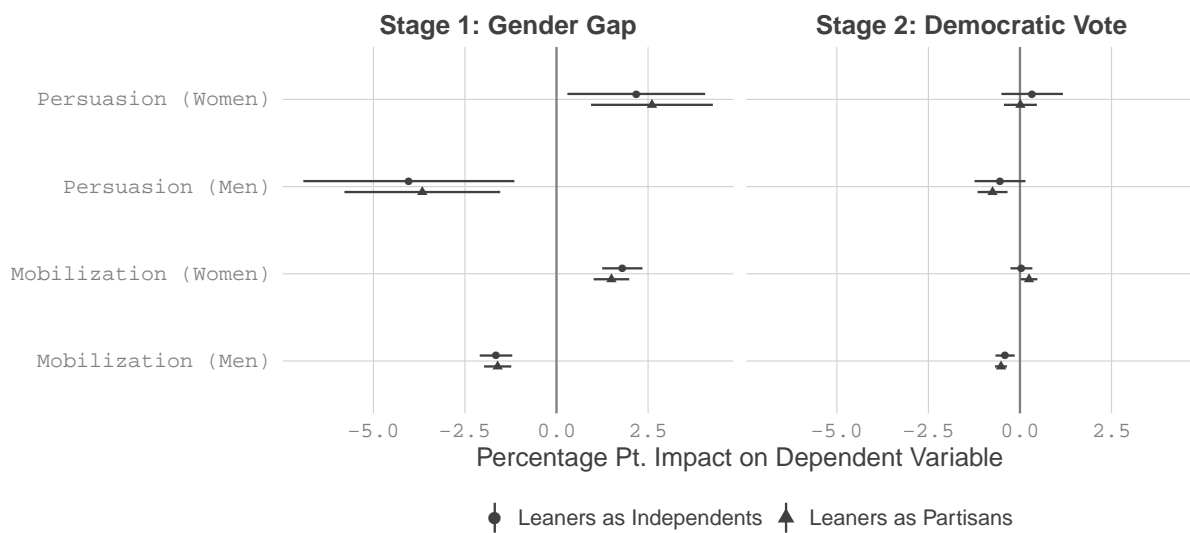
We display the results of each stage in Figure 7. The results from stage one (in the left-side panel) are generally consistent with our expectations of the direction and magnitude of effects. The net mobilization (turnout) and persuasion (defection) measures for women are positively related to the gender gap, while those same measures for men are negatively related to the gender gap. Although the magnitudes of the persuasion variables are not precisely twice that of the mobilization variables, the estimated impact of persuasion is larger than the impact of mobilization on average, and all effects are significant.

The right-side panel in Figure 7 shows the second-stage results.¹⁹ When independent lean-

¹⁹Because each predictor is its observed value multiplied by its stage one coefficient, the second stage predictors are random variables. Standard errors currently do not reflect this reality.

ers are coded as independents, we find that only mobilization among men is a significant predictor of the Democratic vote *via* the gender gap. When leaners are coded as partisans, however, mobilization among both men and women are significant, as is persuasion among men. Though the coefficients from stage two generally reflect our expectations, the coefficients alone do not elucidate the evolving impact of mobilization and persuasion from 1952 to 2012. To tell this story, we now visualize the findings of the two-stage model for each election year.

Figure 7 Results of Two-Stage Model Predicting Gender Gap (Left) and Democratic Vote (Right)




Error bars are 95% confidence intervals

We first visualize how each mechanism impacts the gender gap (stage one) in Figure 8. Lines for men and women are computed by multiplying each mechanism in each year by its estimated coefficient.²⁰ Consistent with the expectations laid out in Table 2, the net Democratic advantage in mobilization (shown the upper and lower left panels) has opposite impacts on the gender gap for women as it does for men. According to our theoretical framework and the stage-one regression, the net number of Democratic voters mobilized from women increases, the gender gap increases *ceteris paribus*; likewise, the gender gap decreases when the number of Democratic voters mobilized from men increases. As the number of Democratic votes mo-

²⁰The lines are equivalent to the values of the predictions inserted into the stage-two regression.

bilized men decreased over time, however (see Figure 6), the negative impact exerted by mobilization among men grew weaker over time. Summing the impact of mobilization among men and women together (in the second column of the figure), we see a gender gap that is consistent with the findings from the observed data in Figure 1—beginning as negative, emerging and then disappearing in the 1970s, and then reappearing permanently in the 1980s and onward.

Figure 8 Impact of Mobilization and Persuasion on the Gender Gap (Stage one)




For persuasion, the roles are flipped. Because persuasion has generally benefited Republican candidates since the 1950s (Shafer and Claggett 1995), we find that persuasion among men has an overall effect of widening the gender gap while persuasion among women decreases it. When we sum persuasion among men and women, we find that persuasion overall had a slightly positive net impact on the gender gap owing to the fact that Republicans netted more votes from Democratic men than from Democratic women in most elections. Persuasion’s impact on the gender gap declines toward zero in recent elections, however, as the prevalence of

swing voting has drastically declined in the era of polarization (Bafumi and Shapiro 2009). As male and female voters have “sorted” themselves into partisan camps, there is less opportunity for campaigns to win over voters from the opposing party (Levendusky 2009).

Now that we have described how each mechanism influences the gender gap over time, we can better interpret the coefficients from the second-stage regressions. Because mobilization among men benefited Democrats before the emergence of the contemporary gender gap (i.e., the variable was positive), its impact on the gender gap was generally negative even though Democrats enjoyed more votes from Democratic men overall. Thus we find that mobilization among men actually increased the Democratic vote before the emergence of the contemporary gender gap. The gender gap emerged in the 1970s and 1980s with a corresponding decrease in the mobilization of Democratic men, so mobilization among men exerts a decreasing impact on the Democratic vote share throughout the time period under examination.

Figure 9 Impact of Mobilization and Persuasion on the Democratic Vote (Stage two)



We estimate that mobilization among women has little effect when leaners are coded as independents (due to a near-zero estimated coefficient in stage two), whereas when leaners are coded as partisans, net Democratic mobilization among women provides a boost to the Democratic vote share of 2 to 4 percentage points. The results of the second stage for men and women combined suggest that mobilization (column 2 in Figure 9) once provided an advantage to the Democrats, but as the number of Democratic male voters decreased from the late 1960s through the 1980s, the gender gap's net influence on the vote (owing to mobilization) shrank toward zero. Following Kaufmann and Petrocik (1999), as men became less Democratic and more Republican, the gender gap widened as the number of Democratic partisan voters in the electorate decreased over time. Upending the folk wisdom, we find that larger gender gaps owing to mobilization have corresponded to a reduction in the number of Democratic voters.

Examining column 4 from Figure 8, persuasion slightly widened the gender gap in the past because the Republican Party collected more votes from male Democrats than the Democratic Party collected from female Republicans. So although net Democratic persuasion among men and women have divergent impacts on the gender gap, the overall balance of persuasion among men and women historically has benefitted Republicans. We see this effect for men and women separately in column 3 of Figure 9—the negative effect of persuasion on the Democratic vote is stronger for men than for women, especially when leaners are coded as partisans. When we sum the effect of persuasion for men and women in column 4, we find that as persuasion's impact on the gender gap has declined in recent elections, so too has its negative net impact on the Democratic vote. This is due to the simple fact that swing voting is far less common in recent decades, and as a result, no party collects many net votes from partisan persuasion in recent elections.

These findings on persuasion also contradict the folk wisdom linking the gender gap and the Democratic vote. Although it is commonly assumed that a reduction in the gender gap should harm the Democrats, the account of persuasion since 1952 finds the opposite to be true. In past decades, persuasion exerted a positive impact on the gender gap, and this impact was

to the net *detriment* of Democratic presidential candidates. As persuasion has declined in the electorate, however, not only does persuasion's contribution to the gender gap also decline, the Democratic Party benefits because fewer Democratic voters are defecting to cast Republican votes.

6 DISCUSSION

There is little support for the assumption that the Democratic Party enjoys a larger vote share when the gender gap is larger. While we find a slightly positive but insignificant correlation between the gender gap and the Democratic vote share, the relationship between these variables is not the result of an unchanging process over time. A careful reexamination of Figure 5 suggests that the positive relationship was evident until the 1980s but has weakened or even reversed since then.

A deeper analysis of mobilization and persuasion also undermines the conventional wisdom. As the mobilization component of the gender gap has stabilized in recent years following the relative increase in mobilization among Republican men, the prevalence of these new Republican male voters has eaten into the Democrats' long-term advantage in partisan mobilization and has decreased the Democratic vote. Similarly, while the persuasion component of the gender gap *used to* contribute to the gender gap in a manner that also hurt Democrats (by boosting Republican votes overall), the disappearance of persuasion in the overall electorate amid partisan polarization has helped Democrats by stemming the level of defection within their ranks of partisan voters.

Because persuasion is worth twice as many votes as mobilization, the confluence of these forces may result in the appearance of an increased Democratic vote. As partisan change drove the gender gap to emerge in the first place, decreases in swing-voting, among Democrats in particular, has shored up the Democrats' base of support. As a result, Democrats retain a greater share of loyal voters on election day even as the number of men within the party has shrunk over

time. In recent elections, although the *gender gap* owes its continued existence to the mobilization of partisan voters, the net influences of mobilization and persuasion on the *Democratic vote* appear largely cancel out one another.

By 2008 and 2012, the only mechanism that continues to exert a non-negligible impact on both the gender gap and the Democratic vote (by way of the gender gap) is the mobilization of women. The fact that this result appears in Figure 9 only when leaners are coded as independents additionally suggests that it is leaning independent women who provide Democrats with the cushion that insulates them from losses in the vote due to partisan change among men. How to interpret this finding depends on one's outlook on independent leaners as partisans or independents. If leaners are independents, these leaners are being *persuaded* to side with the Democratic Party year after year. If leaners are partisans, these women provide an edge in Democratic mobilization that has outlasted the growth of Republican party ID among men.

7 CONCLUSION

Our investigation of the gender gap in postwar U.S. presidential elections has confirmed some popular assumptions while turning up some unexpected relationships. At a high level we dispelled the notion that growth in the gender gap does leads to a bigger vote share for the Democratic candidate. Instead, we contend that this relationship depends on the underlying mechanisms that produce the gender gap in the first place. To understand the mechanisms connecting the gender gap to election results, we developed a general theoretical framework that accounts for both mobilization and persuasion effects of any subpopulations of interest within the electorate.

Applying this framework to the gender gap, we find, contrary to popular wisdom, that while the mobilization of women has had a slightly positive and unchanging effect on the gender gap and the Democratic vote since 1952, it has been swamped by the growing mobilization of Republican men relative to Democratic men. This trend in men's mobilization has nearly elim-

inated the Democratic Party's advantage in mobilized voters, even though it is the largest single contributor to the contemporary gender gap. This runs directly against the popular notion that a large gender gap is good for the Democratic Party. On the contrary, we find that the gender gap in voting emerged due to growing mobilization of Republican men relative to Democratic men. Persuasion has had the opposite effect as mobilization. While the mechanism of mobilization has increased the gender gap over time, the mechanism of persuasion, which exerted a positive impact on the gender gap in past decades, has all but disappeared from the American electorate. Because persuasion's impact on the gender gap served to benefit the Republican Party on balance, the decline in persuasion has increased the Democratic vote even as persuasion's impact on the gender gap has also declined. These findings related to persuasion once again run against the conventional wisdom about the relationship between the gender gap and the Democratic vote share—shrinking the gender gap is widely thought to harm the Democratic Party, but to the extent that declining persuasion has yielded a smaller gender gap, it appears that the net partisan impact has been to protect the Democratic Party from defection in its own ranks.

In short, not only do we find unconvincing evidence for the conventional wisdom in the aggregate, our theoretical framework and two-stage analysis show that there is no static relationship between the gender gap and the Democratic vote. Rather, the relationship between voting gaps and vote outcomes is entirely contingent on differential rates of mobilization and persuasion within each group and across time.

This paper explores what we believe are two important and related themes in contemporary electoral politics. First is understanding how “electoral gaps” between groups in the electorate—in our case, men and women—may lead to an advantage for one party over the other. Second is the more nettlesome task of decomposing these connections into the constituent parts of mobilization and persuasion. These two pieces are the major goals of campaigns and serve as the mechanisms for translating long-term support for the parties into short-term realizations on election day.

There is an irony in finding that mobilization and especially persuasion have mattered less over time and have weakened what might have been a stronger connection between the gender gap and the vote. In fact, a study of only the contemporary era in U.S. politics might find that mobilization and persuasion are modest in size and difficult to distinguish from one another. As a result, recent studies of campaign effects are prone to understate these influences, whether researchers are studying the voting habits of men and women or the voting habits of other groups in the electorate. The sorting of subpopulations into partisan enclaves has dampened the electoral impact of the very mechanisms that created the gender gap in the first place.

REFERENCES

- Axelrod, Robert. 1972. "Where the Votes Come From: An Analysis of Electoral Coalitions, 1952–1968." *American Political Science Review* 66 (1): 11–20.
- Bafumi, Joseph, and Robert Y. Shapiro. 2009. "A New Partisan Voter." *Journal of Politics* 71 (1): 1–24.
- Bartels, Larry M. 2000. "Partisanship and Voting Behavior, 1952-1996." *American Journal of Political Science* (1): 35–50.
- Box-Steffensmeier, Janet M., Suzanna De Boef, and Tse-Min Lin. 2004. "The Dynamics of the Partisan Gender Gap." *American Political Science Review* 98 (3): 515–528.
- Campbell, Angus. 1966. "A Classification of Presidential Elections." In *Elections and the Political Order*, ed. Angus Campbell, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. New York, NY: John Wiley.
- Carroll, Susan J. 2013. "Voting Choices: How and Why the Gender Gap Matters." In *Gender and Elections: Shaping the Future of American Politics*, ed. Susan J. Carroll and Richard L. Fox. New York, NY: Cambridge University Press.
- Chaney, Carole Kennedy, R. Michael Alvarez, and Jonathan Nagler. 1998. "Explaining the Gender Gap in US presidential elections, 1980-1992." *Political Research Quarterly* 51 (2): 311–339.
- Conover, Pamela Johnston. 1988. "Feminists and the Gender Gap." *Journal of Politics* 50 (4): 985–1010.
- Converse, Philip E. 1966. "The Concept of a Normal Vote." In *Elections and the Political Order*, ed. Angus Campbell, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. New York, NY: John Wiley.
- DeNardo, James. 1980. "Turnout and the Vote: The Joke's on the Democrats." *American Political Science Review* 74 (02): 406–420.
- Edlund, Lena, and Rohini Pande. 2002. "Why Have Women Become Left-Wing? The Political Gender Gap and the Decline in Marriage." *Quarterly Journal of Economics* 117 (3): 917–961.
- Hillygus, D Sunshine, and Todd G. Shields. 2014. *The Persuadable Voter: Wedge Issues in Presidential Campaigns*. Princeton, NJ: Princeton University Press.
- Huber, Gregory A., and Kevin Arceneaux. 2007. "Identifying the Persuasive Effects of Presidential Advertising." *American Journal of Political Science* 51 (4): 957–977.
- Kaufmann, Karen M. 2006. "The Gender Gap." *PS: Political Science and Politics* 39 (3): 447–453.
- Kaufmann, Karen M., and John R. Petrocik. 1999. "The Changing Politics of American Men: Understanding the Sources of the Gender Gap." *American Journal of Political Science* 43 (3): 864–887.

- Keith, Bruce E., David B. Magleby, Candice J. Nelson, Elizabeth Orr, Mark C. Westlye, and Raymond E. Wolfinger. 1992. *The Myth of the Independent Voter*. Berkeley, CA: University of California Press.
- Levendusky, Matthew. 2009. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. Chicago, IL: University of Chicago Press.
- Manza, Jeff, and Christopher Uggen. 2006. *Locked Out: Felon Disenfranchisement and American Democracy*. New York, NY: Oxford University Press.
- Manza, Jeff, and Clem Brooks. 1998. "The Gender Gap in US Presidential Elections: When? Why? Implications?" *American Journal of Sociology* 103 (5): 1235–1266.
- Miller, Warren E., and J. Merrill Shanks. 1996. *The New American Voter*. Harvard University Press Cambridge, MA.
- Norrander, Barbara. 1997. "The Independence Gap and the Gender Gap." *Public Opinion Quarterly* 61 (3): 464–476.
- Norrander, Barbara. 1999. "The Evolution of the Gender Gap." *Public Opinion Quarterly* 63 (4): 566–576.
- Norris, Pippa. 2003. "The Gender Gap: Old Challenges, New Approaches." In *Women and American Politics: New Questions, New Directions*, ed. Susan J. Carroll. New York, NY: Oxford University Press.
- Petrocik, John R. 2009. "Measuring Party Support: Leaners are not Independents." *Electoral Studies* 28 (Dec): 562–572.
- Shafer, Byron E., and William J.M. Claggett. 1995. *The Two Majorities: The Issue Context of Modern American Politics*. Baltimore, MD: Johns Hopkins University Press.
- Stanley, Harold W., William T. Bianco, and Richard G. Niemi. 1986. "Partisanship and Group Support over Time: A Multivariate Analysis." *American Political Science Review* 80 (3): 969–976.
- Wirls, Daniel. 1986. "Reinterpreting the Gender Gap." *Public Opinion Quarterly* 50 (3): 316–330.
- Zingher, Joshua N. 2014. "An Analysis of the Changing Social Bases of America's Political Parties: 1952–2008." *Electoral Studies* 35: 272–282.