

The Candidate Supply:
**How the Costs and Benefits of Running for Office Shape
the Democratic Process**

PRELIMINARY AND INCOMPLETE

Draft #1
April 6, 2015

Andrew B. Hall
Harvard University
andrewbenjaminhall.com

Contents

1 Introduction

2 The Candidate-Supply Theory

- 2.1 Demand-Side Politics: A Legacy of the Downsian Model
- 2.2 Theories of Polarization Focus on the Demands of Voters
- 2.3 Why Candidate Positions Are Not Fluid
- 2.4 How The Candidate Supply Matters When Candidate Positions Are Not Fluid
- 2.5 Counterarguments Considered

3 The Nature of Candidate Positions

- 3.1 U.S. House Candidates' Positions: Stable Over Time
- 3.2 Incumbents Do Not Adjust Positions Based On Primary Challenges
- 3.3 Candidate Identity Linked to Candidate Positions
- 3.4 When (If Ever) Do Candidates Change Positions?

4 The Candidate Supply and Constrained Voters

- 4.1 Voters Choose More Moderate Candidates When They Can
- 4.2 The Candidate Supply Constrains Voters

5 Determinants of the Candidate Supply: The Costs and Benefits of Running for Office

- 5.1 How Costs Polarize the Candidate Supply
- 5.2 How Benefits Moderate the Candidate Supply
- 5.3 The Rising Costs of Running for Office
 - 5.3.1 Fundraising: A Growing Burden
 - 5.3.2 Increasing Media Scrutiny of Candidates
- 5.4 The Falling Benefits of Office
 - 5.4.1 Declining Pay for U.S. House Members
 - 5.4.2 Erosion of the Committee System and Diminishing Chances to Influence Policy
- 5.5 Summary

6 Beyond the U.S. House: Considering the Candidate Supply in Other U.S. Elections

6.1 U.S. Senate

6.2 State Legislatures

6.3 Statewide Offices

7 Policy Implications: Altering the Candidate Supply

7.1 Increasing the Benefits of Office

7.1.1 Raising Legislator Salaries

7.2 Decreasing the Costs of Running for Office

7.2.1 Campaign Finance Reform

8 Conclusion

Chapter 1

Introduction

On November 16th, 2013, newly elected Democratic members of Congress sat down to view a presentation by the Democratic Congressional Campaign Committee (DCCC) on how they should allocate their time as first-term incumbents in the U.S. House. On a slide entitled “Model Daily Schedule,” the presenter suggested that new members should plan to dedicate 4 hours per day to “Call Time”—time set aside for making fundraising calls—and another 1 hour per day to “Strategic Outreach”—other forms of in-person fundraising. This is a tremendous amount of time to devote to a single activity, and it is not a pleasant one, either. Reacting to the slides, Congressman John Larson (D, CT) told the *Huffington Post*, “You might as well be putting bamboo shoots under my fingernails.”¹ Describing the off-site call centers that members of Congress use for these activities, Peter DeFazio (D, OR) told *This American Life*: “If you walked in there, you would say, boy, this is about the worst looking, most abusive call center situation I’ve seen in my life.”²

And fundraising is only one of the many burdens members of Congress must bear. Opportunities to legislate—ostensibly the job MCs are sent to Washington to do—are few and far between. In explaining her decision to quit politics, moderate Senator Olympia Snowe (R, ME) pointed to the “dysfunction and political polarization” of Congress.³ The impediments to legislating weigh on the minds of many representatives. Discussing the departures of a number of moderate Senators and members of Congress, Senator Michael Bennett (D, CO) lamented

¹http://www.huffingtonpost.com/2013/01/08/call-time-congressional-fundraising_n_2427291.html, Accessed February 3, 2015. Note that Larson is a long-time member of Congress, not one of the first-year members there to view the presentation.

²<http://www.thisamericanlife.org/radio-archives/episode/461/transcript>, Accessed February 3, 2015.

³http://www.washingtonpost.com/opinions/olympia-snowe-why-im-leaving-the-senate/2012/03/01/gIQApGYZlR_story.html, Accessed February 3, 2015.

that “There are a number of folks who don’t want to come here and participate in the dysfunction.”⁴

To gain the right to bear burdens such as these, would-be members of Congress must first endure an often grueling primary- and general-election campaign. In addition to the tremendous fundraising requirements, these elections place an incredible level of scrutiny on every aspect of candidates’ lives—not only on their careers as politicians and public servants, but on the details of their personal lives and the minute wording of their every utterance to the press. “Electoral politics is a brutal, soul-sucking experience for a candidate”⁵ according to Taylor Griffin, an unsuccessful candidate for the U.S. House in North Carolina’s 3rd district.

The profound costs of running for Congress, and of being in Congress, have not gone unnoticed. Almost daily, it seems, we hear further news about the historic unpopularity of Congress, the gridlock of Congress, and the many burdens of running for Congress. At least four major news outlets have run articles in the past two years with the headline “Why Would Anyone Run For Congress?” or an extremely similar variant.⁶

This book is about who is willing to run for Congress, and, more importantly, about why this question, under-appreciated in the political science literature,⁷ matters for understanding the process of political representation in America. It takes as its starting point the pervading sense—among journalists, politicians, and political observers—that running for Congress has never been more difficult or more costly than it is today, that being in Congress has rarely, if ever, been so frustrating as it is today, and that, as a result, only a particular set of people are willing to enter elections in the first place.

Despite this sense, our pre-eminent political science models of the electoral process focus on the positions strategic candidates choose to take *after* entering

⁴http://www.washingtonpost.com/opinions/bennet-joins-flight-of-the-moderates/2012/02/29/gIQAqhKBjR_story.html, Accessed February 3, 2015.

⁵http://www.realclearpolitics.com/articles/2013/10/24/why_would_anyone_run_for_congress_these_days.html, Accessed February 3, 2015.

⁶These outlets are: RealClearPolitics, The Atlantic, The National Journal, and MSNBC. Found via Google search for the phrase “Why Would Anyone Run For Congress”, Accessed February 3, 2015.

⁷While under-appreciated, it is certainly not a topic without existing research. I review this body of work in Chapter 2.

the race, taking for granted the entry of candidates (Downs 1957). By stressing the fluidity of candidate positions—and demonstrating as a consequence of this fluidity an inexorable movement towards the views of the median voter—these theories are free to ignore the supply of candidates. Who runs for office is not important in the theoretical world in which anyone can adopt the median voter’s views during a campaign.

This focus on fluid positions dominates the literature on legislative polarization in American politics, too. The growth of polarization in U.S. legislatures is well known (McCarty, Poole, and Rosenthal 2006), and many scholars have put forward plausible factors that contribute to this growth. Almost all of these factors, as I will argue, depend on the “demand side” of the political process: the ways in which voters’ preferences produce legislative outcomes of a particular ideological bent by demanding such positions from would-be candidates and only rewarding those who adopt them.

But voters do not always get what they want, and candidates cannot adopt different positions so easily. I argue that candidate positions are in fact more rigid than fluid, and as a result, I focus not just on voter demand but on candidate *supply*. Candidates come to political campaigns with all manner of pre-existing ideological commitments. They have expressed views that they must now hold on to; they have pursued professions that may carry an ideological signal for voters; they have, perhaps, participated in political campaigns, donated to political campaigns, or registered with a party in previous elections, all of which are matters of public record; and they have developed personal commitments to policy positions, commitments strong enough to induce them to run for office in the first place. Though candidates may strategically adopt some positions as they see fit, they are far from unconstrained, and their positions far from fluid.

As a result, the supply of candidates matters. The power of voters in a democracy is fundamentally limited; they are free to choose any of the candidates who run for office, but they cannot elect someone who does not run. Though voters may prefer ideologically moderate politicians, if faced with a choice between only ideological extreme candidates, they cannot send a moderate representative to office, and they cannot turn a more extreme candidate into a more moderate candidate.

The candidate supply thus has the potential to constrain voters and to help determine the ideological positions of policy outcomes in the legislature. What, then, determines the candidate supply? The electoral process, I argue, is akin to a labor market, subject to the laws of supply and demand. Voters demand candidates with certain attributes—e.g., those who hold and offer moderate positions. Would-be candidates supply themselves to voters if they so desire, and if they do, they enjoy benefits both from running for office and, more importantly, from holding office if they win. But, crucially, would-be candidates face costs if they choose to enter the political marketplace. If these costs are too high, or if the benefits of holding office are too low, some would-be candidates will choose not to enter in the first place.

These costs and benefits do not affect all potential candidates equally. More ideologically extreme citizens may be more willing than moderate citizens to bear the costs of running. If this is the case—and I argue in this book that it is—then the supply of candidates will be limited to relatively ideologically extreme ones when costs are high and/or benefits are low. Voters, in turn, will be unable to elect more moderate representatives even if they prefer such candidates. What is more, as the costs of running for office increase, as I argue they have over the past fifty years, the supply of candidates becomes more ideologically extreme.

Departing from the demand-side view of elections, I offer an alternative source of the growing polarization in U.S. legislatures. In this view, it is not the changing preferences of the electorate, nor the changing behavior of primary- or general-election voters that drives polarization—although these factors may well play their part—but rather *it is the changing costs and benefits of running for office*. These changes have been well documented, as the anecdotes at the beginning suggested. Campaigning has become an increasingly burdensome duty, consuming all of the candidate's waking hours. The media's ever-increasing scrutiny on seemingly inane details of candidates' personal lives, not just on their romantic dalliances but, for example, on the contents of college essays written twenty or thirty years ago, contributes all the more to this growing unpleasantness.

At the same time, the benefits of making it into office have diminished, and not necessarily in a way that is equal across ideologies. The last forty years have witnessed a dramatic restructuring of our legislatures (e.g., Hall and Shep-

sle 2014). Where once committees dominated, allowing individual members to participate in the policymaking process through service on particularized committees, party leadership now dominates. Major legislation no longer treads its familiar path through committees, then each chamber, and then the conference committee. Instead, party leaders negotiate directly, “ping-ponging” legislation between the chambers and passing it in identical form with little chance for individual legislators to put in the extra legwork to put their personal stamp on it. Serving in office today seems to be much more akin to being a “foot soldier” for the party, and much less akin to blazing one’s own trail in the legislature.

Together, these changing costs and benefits have drastically shifted the ideological composition of the candidate supply. The costliness of campaigning and the reduced benefits of office mean that, by and large, only relatively ideologically extreme candidates will run for office. And because candidate positions are not nearly as fluid as theories might suppose, the fact that only these candidates will run for office affects the ideological choice-set of voters. *The costs and benefits of office are such that voters are left, in many cases, to elect extremists even though they prefer moderates.*

As well as offering an alternative understanding of electoral politics in the U.S., and a new view of the sources of legislative polarization, this argument also implies a different set of levers for reducing legislative polarization in the future. Rather than—or perhaps, in addition to—reforming the primary system or attempting to create a more informed electorate, reformers should focus on altering the costs and benefits of running for office. Altering these costs and benefits can induce a different and more moderate set of people to run for office, thus giving voters the opportunity to send them to Washington.

I focus on two specific policy instruments, though the candidate-supply theory may support many others, too. First, to increase the benefits of holding office, legislators should be paid far more than they currently are. As I document, legislator salaries, while far exceeding the median American’s income, have fallen relative to other “prestige” occupations. The opportunity costs of becoming a legislator in America are higher than ever. If we want to encourage motivated, educated people of all walks of life to serve in public office, we must compensate them accordingly.

A key advantage to this policy is that it circumvents the usual self-referential problem of policy reform; those whom the new policies affect are precisely those who get to choose whether or not to implement them. Increasing legislator salaries is likely to be one of only a few possible policies that sitting legislators are likely to support. Though the policy must be carefully crafted to focus on encouraging new candidates and not simply on rewarding sitting incumbents without cause—not only because this is sensible policy but because voters surely will not support pointless pay increases to incumbents of historically low popularity—pay raises are uniquely feasible as a policy instrument.

The second policy concerns how campaigns function. Simply increasing the benefits of office will not necessarily improve the candidate supply. Foreseeing the higher benefits, candidates will compete even harder to win office—making the task of running for office even more difficult than it already is, and meaning that the resulting equilibrium may be one in which moderate candidates continue not to run. Thus, raising the benefits in isolation may have little effect. Instead, the increased benefits must be married to a forced reduction in the difficulty of campaigning. I argue that stronger restrictions on fundraising can accomplish this goal. Although campaign finance reform in general has been elusive, such a bill might become feasible if tied to an increase in legislator salaries.

The book is organized as follows. In the next chapter, I offer a theoretical overview of the book’s motivation and arguments. I review the central role the Downsian model has played in studies of elections and of polarization, and I consider how this literature depends on a view of fluid candidate positions. Our existing theories to explain how polarized U.S. legislatures have become all depend on demand-side factors, such as the changing preferences of voters, the changing behavior of primary voters, and/or the geographical sorting of partisan voters into increasingly homogeneous districts. These are all demand-side factors, and they all rely, at least in part, on this concept of fluid candidate positions.

In contrast, I offer theoretical arguments in support of a view of rigid candidate positions, and I discuss the so-called “citizen-candidate” models that motivate my empirical studies. These models focus not just on voter demands but on candidate supply. Inspired by their approach, I propose two key empirical predictions, which I develop in detail. First, as the costs of running for office

increase, the supply of candidates should become more extreme; and second, as the expected benefits of holding office increase, the supply of candidates should become more moderate. I explain these predictions in detail, offer arguments for their plausibility in real electoral environments, and consider possible counterarguments to the candidate-supply theory.

Following this theoretical chapter, I turn to empirical evidence to support my argument. I build this argument up in parts. In the first empirical chapter, I offer new analyses to show how rigid candidate positions in U.S. House elections are. Using a variety of data on candidate positions, I show how candidates rarely change positions over time, and I show how incumbents do not change their positions even when faced with primary challengers from their ideological flank. I also show, instead, how candidate positions *do* seem to vary, systematically, with candidates' underlying identities—consistent with a citizen-candidate type of theory but inconsistent with a purely Downsian one.

Having established the rigidity of candidate positions, in the next empirical chapter I turn to the preferences of voters, showing that they display a marked preference for more moderate candidates, *among the supply of possible candidates*. Combining primary- and general-election data for the U.S. House, I show that the “as-if” random assignment of a more ideologically extreme candidate causes a significant electoral penalty to the party, on average. If voters prefer moderates, why do we observe so much polarization in the legislature? Because candidate positions are rigid, if the candidate supply is relatively extreme then so, too, will be representatives. And this is indeed the pattern I uncover. Using data on the positions of U.S. House candidates—both winners and losers—I show how the candidate supply has become more extreme over time. The growth in legislative polarization is mirrored by a growth in candidate-supply polarization.

Where has this growth come from? In the next chapter, I examine variation in the costs and benefits of running for office, offering empirical support for my theoretical arguments. The supply of candidates becomes more extreme when the costs of running for office are higher and when the expected benefits of holding office are lower. To establish this first claim concerning costs, I examine the ideological positions of state legislators running for the U.S. House in cases when they have to give up their state legislative seat vs. in cases when they can run

without risking their current seat. Ideologically moderate candidates are much more likely to run when this cost is lower than when it is high, while ideologically extreme candidates are equally likely to run regardless of the cost. To establish this second claim regarding benefits, I look at the effects of incumbency on the candidate supply. Incumbency status represents a downward shift in the expected benefits of holding office for a would-be candidate because it lowers the probability that she will win office if she runs. Consistent with expectations, the candidate supply is more ideologically extreme when an incumbent holds office than when there is an open seat.

These analyses show how the costs and benefits of running for office affect the candidate supply. Is it true, as I have supposed, that the costs have gone up and the benefits down over the past forty years, as polarization has increased? In the final part of the chapter, I present a variety of statistical analyses to show how much more costly it has become to run for office, and to suggest that, at the same time, the benefits of holding office have likewise decreased.

Having built up these empirical supports, I turn to potential policy implications. I discuss possible reforms that would decrease the costs of running for office and increase the benefits of holding office. Reforms to campaign finance that reduce the amount of time candidates have to spend raising money seem especially promising for reducing the costs of running for office; increasing legislator pay may be one lever for increasing benefits, although it must be carefully weighed against other drawbacks. I also survey the use of such reforms in other electoral democracies, and I use these comparative cases to fine-tune possible reforms for the U.S. context.

In the final chapter, I conclude by discussing the implications of my argument for our electoral theories. Focusing on demand-side electoral politics has produced incredibly useful discoveries about American politics. But when we focus on supply and demand together—taking the candidate supply as a serious constraint to electoral outcomes—we uncover different and important implications for elections and representation.

All representative democracies face a fundamental problem of candidate supply. Running for office requires candidates to pay a variety of personal costs. The set of potential candidates who would be “best” for society are not necessarily

those willing to pay these private costs—indeed, in many cases they are unlikely to be. This book is a first step in understanding the factors that exacerbate or ameliorate this problem in U.S. elections.

Chapter 2

The Candidate-Supply Theory

Why do candidates adopt the positions that they do? This is a fundamental question in the study of democracies. In the American context, as in many others, elections are explicitly intended to force candidates to cater to the needs of citizens. In Federalist 52, Publius (1787) famously declares:

As it is essential to liberty that the government in general should have a common interest with the people, so it is particularly essential that the branch of it under consideration should have an immediate dependence on, and an intimate sympathy with, the people. Frequent elections are unquestionably the only policy by which this dependence and sympathy can be effectually secured.

Do elections “secure” this “dependence”? To answer this question, and to understand how democracy functions, we must understand how and why candidates come to offer the positions that they do.

What is more, only by doing so can we attempt to explain why, in recent years, the positions of candidates of the two major parties have diverged so dramatically. Polarization in U.S. legislatures is near all-time highs (McCarty, Poole, and Rosenthal 2006), and much of this polarization is the result of the widely diverging positions that Democratic and Republican candidates offer to the same set of constituents—and the same median voter—in each Congressional district (McCarty, Poole, and Rosenthal 2009). To explain this pattern, too, we must understand the roots of candidate positions in U.S. elections.

In this chapter, I review the dominant views of elections and polarization in the American Politics literature, and I explain how these views are united in an almost exclusive focus on the demands of voters, rather than the supply of candidates. The focus on demand-side politics has been driven, I argue, by models that assume candidates’ positions are fluid. With fluid positions, the

candidate supply places little constraint on voters so long as a minimum number of candidates are willing to run. On the other hand, if candidate positions are rigid, as I document in the next chapter, then we miss an important part of the political process, one which helps explain polarization, by focusing only on the demands of voters and not also on the supply of candidates.

2.1 Demand-Side Politics: A Legacy of the Downsian Model


The predominant model of candidate positions, one that provides the foundation for thinking about democratic elections, is found in Downs (1957). Cited more than 20,000 times,¹ Downs' book offers a clear theory that predicts that candidates will *converge*, offering the positions that the median voter prefers in order to secure a majority of votes.

In its simplest form, the model that Downs develops considers two candidates running to represent a constituency. The ideological preferences of each voter can be arrayed along a single line, and each candidate chooses a platform that is represented as a point along this line. Voters support whichever candidate offers a platform closer to their own preferred point. Logically, the only inevitable position for candidates to offer is that of the median; if either candidate offers any other position, either to the left or to the right of the median, then the other candidate can offer the median position and win with certainty.

The Downsian model is a “demand-side” model of the political process. That is to say, the model is about the demands of voters, taking as given the supply of candidates. It represents candidates' positions as entirely fluid; they are free to adopt any positions they want in order to gain election. Any candidate in the race can, for example, adopt extreme views or, as the model predicts will occur in equilibrium, can adopt the views of the median voter. Candidates choose positions strategically in response to each other's behavior and the demands of their voters. As a result, the candidate supply matters very little; so long

¹This estimate comes from Google Scholar, accessed February 3rd, 2015.

Figure 2.1 – The Candidate Supply is Unimportant in the Downsian Model. Regardless of personal preferences, Downsian candidates move to the median based on voter demands.



as candidates exist—which the model assumes—their identities do not matter. Whoever they are, they can adopt the median voter’s preferred positions

Figure 2.1 explains this graphically. The top part of the plot represents the ideological distribution of the candidate supply. The bottom part depicts the same distribution again, for voters. The two distributions are identical in the plot because any voter can, in theory, become a candidate if she desires. The black dot in the lower plot represents the median voter. Suppose a Democrat and a Republican with left and right wing views, respectively, are running for office. Regardless of their personal views, they converge to the median voter’s position, as the arrows represent.

This way of thinking about elections has guided a tremendous amount of research in American politics, and for good reason. It provides a clear set of assumptions that lay bare why candidates may, in many cases, “move to the middle.” Do candidates of the two parties converge to the positions of the median voter? Scholars who have studied this question have found that the answer is “no.” Examining the positions of U.S. House candidates over a long time period, Ansolabehere, Snyder, and Stewart (2001) finds a marked gap in positions

between Democrats and Republicans even when running for election in districts with the same underlying partisanship. Going further, Lee, Moretti, and Butler (2004) uses a regression discontinuity design to focus on districts where a Democratic or Republican candidate is “as-if” randomly assigned to hold office. Districts that quasi-randomly receive a Democrat see much more liberal roll-call voting from their representative than do the districts that receive the Republican incumbent, even though the underlying preferences of the two sets of districts are, on average, identical.

While these findings raise questions about the Downsian model’s ability to predict dynamics in real U.S. elections, they also highlight its value as a generator of hypotheses. By asking why empirical results do not match the Downsian model, we can think about which of the model’s assumptions are likely not to be met, in practice. Many papers have contributed to the literature in this manner.² In this direction, I argue that the failure of the median voter theorem is related to a broader phenomenon in U.S. politics. In particular, if candidates do *not* have fluid positions, and if the supply of candidates is constrained, then voters may not be able elect a candidate offering median-like positions. After describing in the next section how this view differs from most work in the polarization literature, I lay out my argument for why the candidate supply is important.

2.2 Theories of Polarization Focus on the Demands of Voters

Like many electoral literatures in political science, the extensive literature on legislative polarization in American politics takes its cue from the Downsian model. The main explanations for the large and growing polarization of U.S. legislatures rely on demand-side factors: the changing preferences of voters; the changing behavior of primary voters; and/or the changing distribution of voters across districts.

²For a comprehensive review of this literature up to 2004, see Grofman (2004). For a comparative examination of centripetal and centrifugal forces in elections in general, see Cox (1990). Different from the candidate supply factor I consider, Palfrey (1984) offers a theory of third-party candidate entry to explain divergence.

Perhaps the defining work in the literature on polarization is McCarty, Poole, and Rosenthal (2006). McCarty, Poole, and Rosenthal's book presents roll-call based evidence for the sharp growth in legislative polarization since the 1970s and identifies the growth in U.S. income inequality—along with a concurrent change in the way the income distribution and changing voter preferences have altered the strategic decisions of MCs—as a key factor in this growth.

A large body of other work enumerates additional changes in the preferences and behaviors of citizens that may also influence polarization. Abramowitz (2011), for example, focuses on the differences in—and the changes over time of—voter engagement across the ideological spectrum. Bishop (2009) discusses the sorting of citizens across districts as another possible factor (though see Abrams and Fiorina (2012) for counterarguments). Levendusky (2009) studies a different kind of sorting—the sorting of liberals and conservatives into the Democratic and Republican parties, respectively, which he argues is responsible for the changing nature of polarized politics in the U.S.

Relatedly, a large literature points to the U.S. system of primary elections, beginning with a theoretical literature on two-stage elections (Aranson and Ordeshook 1972; Coleman 1971; Owen and Grofman 2006). Burden (2001) points to primaries to explain the failure of U.S. candidates to converge to the median voter. Brady, Han, and Pope (2007: abstract) likewise argues that “primaries pull candidates away from median district preferences.” In perhaps the strongest version of these claims, Pildes (2011) proposes abolishing the primary system because of its hypothesized effects on polarization. On the other hand, Hirano et al. (2010) examines changes in polarization as states implemented primary elections and finds no evidence of a polarizing effect. McGhee et al. (2014) investigates the effects of changes in the type of primary election (open vs. closed) and again finds little relationship with polarization. Whether primary elections play any role in legislative polarization thus remains an open question.

Finally, a separate literature focuses on redistricting as a causal factor, hypothesizing that the creation of more partisan districts, i.e., districts with a greater proportion of voters that support one party or the other, induces legislators to take more extreme positions. Carson et al. (2007: abstract) reports “modest” effects of redistricting on polarization and concludes that “redistricting

is one among other factors that produce party polarization in the House.” Theriault (2008) makes similar arguments, as do a number of journalistic accounts of the process (see for example Eilperin 2007). However, McCarty, Poole, and Rosenthal (2009) presents clear evidence against the hypothesis.³ First, as they point out, polarization appears to be just as high in the U.S. Senate as in the House, despite the fact that redistricting only occurs in House elections. Second, they also show that most of the polarization in the House actually comes from *intradistrict* polarization, the degree to which the two parties diverge within a given district, rather than from *interdistrict* polarization, as would be caused by gerrymandering.

These ongoing literatures are united in their focus on the demand side of politics. In all cases, the key explanatory factor is the preferences of voters, translated into polarization directly through the way they vote in either the primary or general election. Whether or not voters themselves have polarized over time (see Fiorina, Abrams, and Pope (2005) for arguments to the contrary), the preferences of voters need not be the only input into polarization. The behavior of candidates, separate from voters—the “supply” side—is likely to play just as important a role. Fiorina, Abrams, and Pope (2008) raise this point in their discussion of the literature. They write: “When statistical relationships change, students of voting behavior have a tendency to locate the source of the change in voter attitudes, but unchanging voters may simply be responding to changes in candidate strategy and behavior (556).”

By focusing on the demand side, these literatures often also rely on an implied model of fluid candidate positions. The behavioral literature on polarizing voter preferences implicitly assumes that such preferences map into the legislature via strategic candidates who cater to voter demands. If candidates are unwilling or unable to change their positions, then shifts in the electorate’s preferences will not necessarily appear in the legislature. An alternative mechanism would be one of incumbent *replacement*, in which the shifting preferences of voters bring new candidates offering new positions into Congress. Distinguishing replacement from so-called candidate “adaptation” is the subject of an extensive literature which comes to mixed conclusions (Theriault 2006). Though candidate rigidity

³Also see Masket, Winburn, and Wright (2012).

is likely to make adaption an unlikely mechanism—as evidence in Chapter 3 of this book suggests—either story remains one of demand-side forces.

Likewise, for primary electorates to pull candidates away from the district median, candidates must be able to *change* their positions in response to primary voters' demands. Strategic incumbents may have incentives to ward off primary challenges by hewing closer to the party's stances (and farther from her district's median voter's views), but doing so in practice requires deviating from one's chosen positions, with all the personal and strategic costs that that entails.

Finally, explanations that rest on changing voter support for differing candidate positions across districts and over time assume a steady supply of those positions amongst the candidate pool—either because such candidates are always willing to run or, more likely, because of an implicit model of fluid positions in which candidates can adapt to the needs of their electorate at any time.

Such a focus is not inappropriate; the demands of voters are a key driving force in democratic politics. However, they are only half the story. Like in any economy, demand must be met by supply. Since our system of government requires that representatives enact policy on behalf of voters, voters must find candidates willing and able to translate their views into public policy. When they cannot do so, the candidate supply becomes a binding constraint on their objectives. If candidate positions are not fluid, and if not all types of candidates are willing to run for office, then this constraint can become binding—a fact that demand-side studies of elections and polarization cannot, in isolation, address.

2.3 Why Candidate Positions Are Not Fluid

In this book, I argue that candidate positions in U.S. elections are not nearly as fluid as these literatures suppose. Later in the book, I present a bevy of evidence on candidate positions in U.S. House elections to support this argument. Here, I lay out reasons for *why* candidate positions may be rigid.

First, candidates may hold strong personal preferences over policy—indeed, this is likely part of why they are willing to run for office in the first place, as I argue. Many theories allow for at least some room for personal preferences in candidate positions and behavior (e.g., Aldrich and Rohde 2001; Cox and

McCubbins 2007; Wittman 1983). Articulating this view especially strongly, Wittman (1977: 180) writes: “candidates view winning as a means to policy.” Personal conversations with congressional staffers reaffirm this view; staffers I have spoken to are often struck by the deep commitment, or even obsession, their bosses have to particular policy issues. Though re-election concerns are never far from members’ minds, neither are the specific policy goals that, in many cases, drove them to seek office. Formally interviewing a variety of congressional staff, Jacobs and Shapiro (1997: 3) concludes that “the policy positions of members were guided by their personal beliefs, ideology, and judgments.”

Caring about policy does not guarantee that candidates will not change their views. The need to gain election in order to implement policy may still force candidates to offer views to voters other than their own. But it makes it costly for them to do so. If these costs are high enough, candidates may not be willing, in all cases, to alter their views for strategic gains. This is especially true when candidates are uncertain about voter preferences. Probabilistic voting models—models that incorporate uncertainty about the location of the median voter—typically predict that candidates with personal preferences will move towards their own views and away from the middle (Calvert 1985; Wittman 1977). Empirically, we know that U.S. House candidates have little information to go on when determining positions. As Jacobs and Shapiro (1997: 3) reports, “one reason members do not rely on opinion polls is that they do not trust them.” Possessing personal views, when combined with tremendous uncertainty about the most politically advantageous positions, provide a powerful incentive to candidates to offer positions similar to their personally favored ones.

Second, voters may punish candidates for changing positions, forcing candidates to hold fixed positions over time. Tomz and Van Houweling (2015) presents experimental evidence that voters do not react favorably to “flip-flopping,” interpreting changes in candidate positions to be a sign of weak character. After showing survey takers “policy histories” of hypothetical candidates, the authors asked respondents to rate them on the basis of five traits labeled “Honest,” “Moral,” “Strong Leader,” “Knowledgeable,” and “Open-Minded.” When respondents saw candidates with varying policy histories, they rated them consistently lower on all four traits other than “Open-Minded.” As the authors write in a detailed

synopsis of the experiment, “In general...voters drew negative inferences about the character of candidates who changed positions over time.”⁴ The Tomz and Van Houweling (2015) results are experimental, but they are likely to translate into the world of real elections. Punishment for flip-flopping need not require voters to recall previous candidate positions; there is no shortage of interest groups, or opposing candidates, ready to point out such behavior when it occurs.

Focusing on senatorial elections—but with obvious applicability to House elections, too—DeBacker (2015) shows how incumbents are hurt by flip-flopping. Investigating changes in senator roll-call positions over time, the paper estimates “significant costs of of changing position” (109). Closely related to this book’s arguments, DeBacker uses the results to discuss the possibility that incumbent rigidity is, at least partially, responsible for the observed failure of candidates to converge to the median voter.

Third, and finally, candidates may simply be unable to convince voters that they are being honest if they change positions—even if voters would not punish them for these changes if they could believe them. This problem of commitment is at the heart of citizen-candidate models, which I review in the next section. One way to rationalize the problem is that voters have external information about candidates’ backgrounds, previous statements, previous political behavior, and more. Voters thus have little reason to believe a candidate if she claims to have a new position on an old issue, or if she offers a position on a novel issue that seems at odds with her established positions on similar, existing issues.

For these and other reasons, candidate positions are difficult to change, and candidates will change them only rarely. Chapter 3 presents evidence consistent with this hypothesis. Before we turn to that empirical analysis, I consider what rigid candidate positions imply for our theories of candidate positioning and elections.

⁴http://web.stanford.edu/~tomz/working/TomzVanHouweling_PoliticalRepositioning.pdf, p.7, Accessed February 6, 2015.

2.4 How The Candidate Supply Matters When Candidate Positions Are Not Fluid

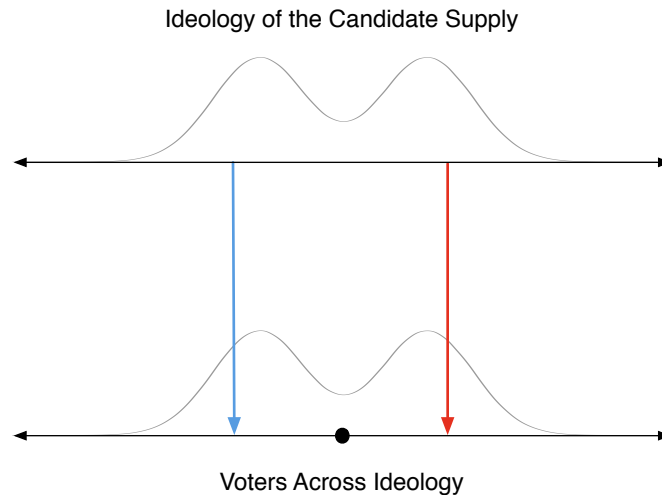
If candidate positions are rigid and not fluid, then the identities of those willing to run—and what kinds of positions those people hold—matter. We therefore need to understand the supply of candidates; we need to understand who is willing to run for office. In this section, I describe a simple theory of the candidate supply based closely on previous theoretical and empirical work.

In keeping with the discussion above, I take as my starting point the view that candidates are motivated not just by the desire to hold office but also by personal policy preferences. Though candidates have strategic incentives to offer positions that the median voter likes, they are, by and large, unwilling and unable to do so. Elections are therefore contests between those willing to enter the race—candidates whom voters choose between on the basis of their ideological views and other characteristics.

Though my arguments are largely empirical—and will be laid out in subsequent chapters of this book—they are motivated by a set of theoretical models known as “citizen-candidate” models (Besley and Coate 1997; Osborne and Slivinski 1996). These models preserve the Downsian view of simple, single-dimensional voter preferences, but they depart radically from the idea of fluid platforms that can adapt to meet the desires of the median voter. Rather, in citizen-candidate models, candidates can only offer their true, personally preferred policy positions, regardless of where these lie. The median voter can thus enter the race and win, in many cases, due to her advantaged position in the middle of the spectrum, but if she does not enter then voters simply will not have a choice of the median platform. Figure 2.2 depicts this relationship between the candidate supply and realized candidates. Unlike in Figure 2.1 from before, candidates who enter now must offer their own views to voters. Thus, if non-median candidates enter, voters must choose among them.

Because citizens can only offer their preferred views, and not simply whatever the median voter wants, who chooses to run for office matters in citizen-candidate models—matters in the sense of affecting the final policy outcomes that voters

Figure 2.2 – The Candidate Supply Matters When Candidates Can Only Offer Their Own Views. Unlike in the Downsian model, in the citizen-candidate model the identities of the candidates who enter determines the ideological choices that voters face.



receive. It thus differs from the Downsian model by modeling simultaneously the electoral process, conditional on who runs, and also the decision to run in the first place. As Besley and Coate (1997: abstract) writes, “...it is a conceptualization of a pure form of representative democracy in which government is *by*, as well as *of*, the people” [emphasis original].

In such models, citizens’ utility depends on two separate factors: first, how far realized policy is, ideologically, from their personally preferred position; and second, the net benefits of running for office (the expected benefits of winning minus the costs of running), if they choose to enter the race. Though any candidate can choose to enter, she must pay a cost for doing so. If a candidate wins election, she receives a fixed, non-ideological benefit in addition to the opportunity to implement her preferred policies. The decision to enter thus depends on: the size of the benefit for winning; the probability of winning the election; the size of the cost for running; and, crucially, her views relative to those of her potential opponents. If she thinks a candidate with similar views to her own is running, then she may be willing to sit out and avoid paying the cost of running, “free-riding” off the candidacy of her like-minded fellow citizen.

Figure 2.3 – Decision Entry of a Potential Moderate Candidate.


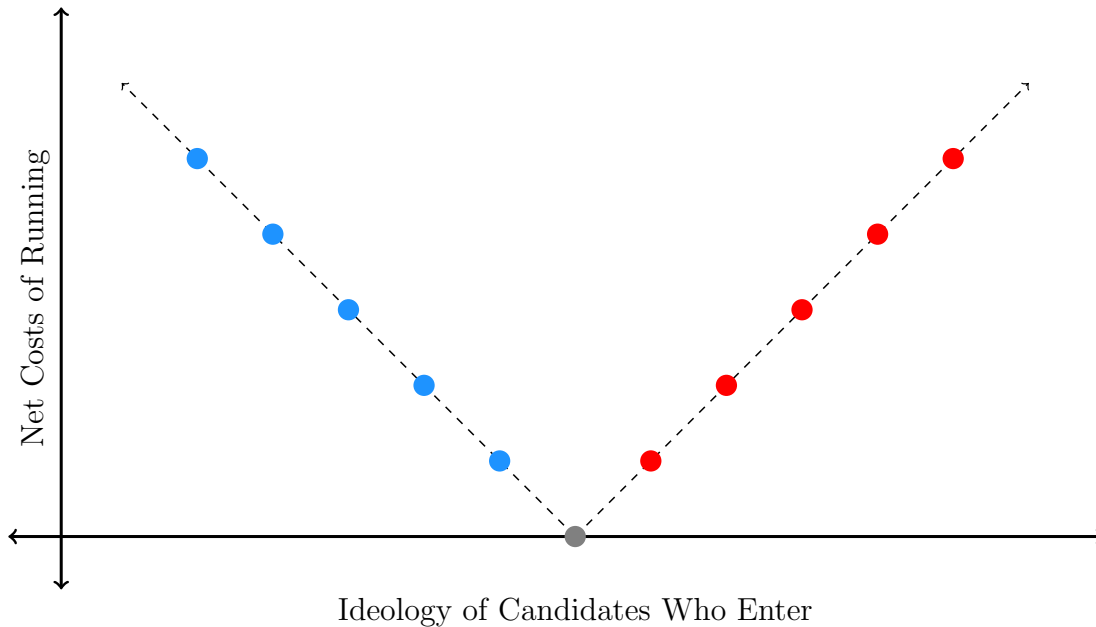


Figure 2.3 helps explain the decision from the point of view of a moderate potential candidate. The figure presents a hypothetical election in which two candidates—one to the left and one to the right—have already decided to run for office. The middle point represents a potential moderate candidate deciding whether or not to run. *If* she runs, she wins the election because she is located at the median, while the other two candidates each accrue only a relatively small portion of more extreme votes. However, to run she must pay the costs. What happens if she does not run? If she sits out the race, she sees either the election of a relatively far left or relatively far right candidate.

But now consider the situation from the point of view of the “Left Extremist” candidate. If *she* bows out while the moderate is not running, the far right candidate wins, representing policy positions twice as far from her as from the moderate candidate. And this situation is the same, in reverse, for the “Right Extremist” candidate. In this case, *the extreme candidates face worse consequences from not running than does the moderate candidate*. This is why costs of running may deter moderate candidates but not extreme candidates, even if both types of candidates bear the same underlying costs—that is to say, even if extremist candidates are just as sensitive, in their underlying preferences, to the non-ideological costs of running as are moderates. Even in this simple model, the fear of giving control to the other side leads more extreme candidates to be more willing to run for office than more moderate candidates.

Figure 2.4 – Costs of Running for Office and Ideological Divergence of Candidates. For a realistic range of the net costs of running for office, higher costs lead to candidates with more distinct positions.



This becomes all the more true as the costs of running increase. The higher the costs of running are, within limits, the more extreme the candidates in the race will be. Only candidates farther to the left or farther to the right will fear the outcome if they do not run enough to be willing to run when the costs are high, and only candidates that extreme—and no more extreme—will remain sufficiently close to citizens in the middle (including the median) to deter them from entering the race.

Figure 2.4 helps make this clear. On the horizontal axis, the plot depicts the ideology of candidates who enter the race. The vertical axis represents the costs of running for office. When costs are zero, the median citizen enters the race and faces no opposition (since she wins for sure against any other single candidate). As the costs increase, distinct candidates on either side of the median enter the race, and the median herself does not. The higher the costs are, the wider is the gap between the candidates who enter. Of course, this relationship does not hold to infinity; if costs are sufficiently high, other equilibria obtain (for example,

if costs are especially high, it is possible that no one will enter). However, for a reasonable range of costs—those consistent with the fact that we always do see candidates for U.S. House elections—the model contains equilibria of this form. As the net costs rise, i.e., as the costs of running go up or the expected benefits of holding office go down, the candidates who enter will be farther apart, ideologically.

The citizen-candidate model is highly abstract, and I do not claim that it fits reality perfectly. Among the important factors it excludes are: primary elections, which affect entry decisions; candidate attributes other than ideology, which affect electoral outcomes; and differential costs and benefits across candidates, which will affect equilibrium decisions and outcomes. Nevertheless, it isolates a key feature of the U.S. electoral process. The decision to run for office is a costly one, and one that is surely made with a gimlet eye towards the available benefits should one win election. These decisions, and the factors that go into them, are not fixed across the ideological spectrum. Candidates farther from the center have more to gain from the opportunity to implement their preferred policies, and more to lose from foregoing the opportunity to do so, than do candidates whose views are near the center.

This feature of the model accords with the ways potential candidates discuss the decision to run for office. Matt Salmon (R, AZ), for example, returned to Congress after a long absence citing a “concern for future generations.” He told *RealClearPolitics* that he “came back out of fear.”⁵ Describing the mindset of many potential candidates dismayed at the current state of policy, the article concludes: “the fundamental way out of the current condition remains a straightforward one: grab the reins yourself.”

This theory yields two main testable hypotheses. First, if candidates behave in this manner, then *as the costs of running for office increase, the supply of candidates should become more extreme*. As running becomes more costly, citizens towards the middle of the spectrum become increasingly unwilling to run.⁶

⁵http://www.realclearpolitics.com/articles/2013/10/24/why_would_anyone_run_for_congress_these_days-2.html, Accessed February 9, 2015.

⁶In citizen-candidate models, it is possible for costs to be so high that no one runs for office. I do not consider this situation since, empirically, we always observe at least one candidate running for U.S. House seats.

Relative to moderate citizens, more extreme ones are more willing to keep running because they fear the outcome if they do not run more. Put another way, *moderates should be more sensitive to the costs of running for office*. Second, by the same logic, moderates should also be more sensitive to the benefits of holding office. In particular, *as the expected benefits of office increase, the supply of candidates should contain more moderates*. These predictions become stronger, too, if moderates bear higher non-ideological costs of running for office than do extremists.

In focusing on the potential costs and benefits of running for office, I follow a longstanding literature on “candidate ambition,” stemming from Schlesinger (1966) and Black (1972). In these works, would-be candidates are thought to make similar cost-benefit analyses when deciding whether (and when) to run. A closely related body of work, building off of these works, examines how incumbents “scare off” quality challengers (Cox and Katz 1996; Jacobson 1989; Levitt and Wolfram 1997). These literatures have focused primarily on other attributes of potential candidates—like quality or sex—but the logic applies with equal force to ideology.

Separate from the ambition literature, the candidate-supply theory is also motivated by evidence presented in Thomsen (2014), which establishes that more extreme Democratic and Republican state legislators are more likely to run for the U.S. House than are their more moderate colleagues. Thomsen (2014) and Thomsen (N.d.) focus on how would-be candidates consider their degree of “party fit” in the legislature before deciding whether to run. Party fit is a possible benefit (or cost) of holding office. As Thomsen (N.d.: 6) explains: “The central claim is that candidate ideology—and more specifically, the congruence between a candidates ideology and the ideological reputation of her party—influences the decision to run for office.” Though I do not test the party-fit hypothesis directly—I focus instead on other sources of variation in the costs and benefits of office—it is likely to play an important role in determining the costs and benefits of running for and holding office.

2.5 Counterarguments Considered

Before proceeding to the empirical analyses, I pause to consider important alternatives and counterarguments to the candidate supply theory I have outlined above.

Non-Ideological Candidate Attributes

The candidate-supply theory, as I have sketched it out, concerns only ideology. It does not consider other attributes of those who choose to run for office and those who do not. Voters care about many things, not just the ideological positions that candidates offer. These other things include concepts like “quality,” “character,” and “honesty.” We know that quality, especially, exerts a separate force on the electoral fortunes of candidates (Jacobson 1989, 2012).

Nevertheless, there is a strong consensus that the ideology of our representatives is an important and vexing issue in U.S. politics today. Though we must keep in mind the many factors left out of the candidate-supply theory, it, like any theory, gains tractability and clarity through its more narrow focus. The candidate supply is an important constraint in the democratic process. While I document this constraint only in terms of ideology in this book, I hope that future work will extend the principles of the theory to these other traits, too.

Mutidimensional Voter Preferences

Even with this singular focus on ideology, other issues may persist. Recent work documents how voter preferences may not be nearly as simple as in the Downsian model. Broockman (2014a) presents the results of a survey that suggests voters hold relatively extreme views on many policy issues. When the answers to these policy questions are averaged together, voters wind up looking “moderate” but only because the averaging collapses over a variety of dimensions on which respondents hold extreme views in conflicting directions. Broockman thus concludes that the story often told in the U.S. elections literature, one in which many voters are stuck in between the positions of their candidates (Bafumi and Herron 2010), is not the whole story.

While important for other reasons, this argument is not problematic for the candidate-supply theory because the theory concerns voter preferences over candidates and not issues. In Chapter 4, I document that, given the choice of moderate or extreme candidates, voters in actual elections consistently choose more moderate candidates. This is not to say that the underlying issue-specific preferences of voters are unimportant, but rather to say that their preferences over candidates may differ from their issue-specific preferences. It is important to consider what voters want, issue by issue, when considering the *normative* implications of our process of representation, but to understand the choices they face among *candidates*, it is logical to focus instead on how voters go about choosing among candidates.

Party Activists

Another burgeoning literature casts doubt on the very notion of modeling elections as a simple interaction between candidates and voters. Bawn et al. (2012) highlights the role that elites and interest groups play in recruiting candidates, providing necessary support networks to make chosen candidates viable, and in determining the issues that dominate the campaign process. In the same vein, Noel (2014) argues that the issues that arise in the legislature arise first among elite actors outside the legislature. Polarization is thus one product of shifting elite preferences in society, not a product of electoral forces *per se*.

Though quite different in its focus on party networks, it is possible to reconcile this literature and the candidate-supply theory. Party networks exist in part to influence the costs and benefits of running for office—inducing the candidates that party activists like to run while discouraging those that they do not. This is very likely another source of the costs that moderate candidates face when considering a run for office. Indeed, we know that candidate recruiters report preferring non-moderate candidates in surveys (Broockman et al. 2014), and thus are more likely to encourage non-moderate candidates to run. In this book I focus on other sources of variation in the costs and benefits of running for office, but the literature on party activists and party networks provides another promising source of such variation.

Summary

In this chapter, I have reviewed the Downsian model and its contributions to the study of polarization, and I have argued that the majority of work on the subject focuses on a demand-side view of the political process, taking for granted a steady supply of candidates willing and able to offer any ideological positions that voters desire. While the factors that this literature has identified are no doubt important, they offer only an incomplete view of the electoral process. Understanding this process requires examining how the preferences of voters *interact*, in a strategic environment, with candidates. Part of this interaction is the decision of citizens to offer themselves as candidates to voters.

Accordingly, I have put forward a theory that focuses on the candidate supply. Not all citizens will run for office at all times; as a result, their choices may constrain voters, consigning them to elect candidates with positions other than those they most desire. In particular, I have described a theory that explains how more moderate candidates will be more sensitive to the costs of running for office—and to the benefits of holding office—and that thus predicts that the candidate supply will be more ideologically extreme when these costs are higher, and/or when these benefits are lower. The candidate-supply theory links the electoral demands of voters to the supply of willing candidates, and in so doing, it helps explain why polarization has risen so much in our legislatures.

I now turn to a series of empirical analyses. First, I validate the underlying assumptions of the candidate supply theory. Following that, I carry out tests of its implications concerning changes in the candidate supply due to changes in the costs and benefits of running for office.

Chapter 3

The Nature of Candidate Positions

The first key link in this book’s argument is that candidate positions are, by and large, *rigid*. Rather than moving fluidly in response to the demands of voters, candidate positions move only haltingly if at all. As a result, who is willing to run for office—and the views that such people hold—are an under-appreciated source of the positions legislators take in U.S. legislatures.

In this chapter, I review existing empirical evidence on the nature of candidate positions, and I offer a variety of analyses to support the rigidity argument. Using estimates of candidate ideology based on campaign contributions, I show that candidates tend not to change their positions over time. I also explore how incumbents react to primary challenges by extremists, again documenting ideological rigidity. Even when challenged from their parties’ flanks, incumbents do not change their positions, on average. Finally, I show that the underlying identities of candidates are associated with their positions, suggesting that candidates choose positions for personal reasons and not purely for strategic ones like in the Downsian framework. Taken together, the evidence in this chapter establishes why understanding the candidate supply is so important. Candidates emerge committed to a variety of ideological positions that will not change, regardless of strategic context or the preferences of the electorate. Who emerges thus matters for voters.

3.1 U.S. House Candidates’ Positions: Stable Over Time

Perhaps the best existing evidence for the candidate rigidity hypothesis comes from Poole and Rosenthal (2000). Examining how U.S. House legislators’ posi-

tions, measured via their roll-call votes, change over time, they conclude: “we find remarkable and increasing stability...Members of Congress come to Washington with a staked-out position on the continuum, and then, largely die ‘with their ideological boots on’” (8).

The evidence presented in Poole and Rosenthal (2000) concerns U.S. House incumbents only. While suggestive, it is possible that the stability they detect is an artifact of the manner in which legislators, and especially the majority party, control the roll-call voting agenda. In addition, regardless of these agenda effects, the results only speak to incumbents, and thus do not indicate whether the entire candidate supply—including challengers as well as open-seat candidates—displays the same ideological rigidity.

Here I extend the analysis to U.S. House candidates, not just incumbents, using the donation-based measures of candidate ideology from Bonica (2014), called CFScores, downloaded online from the Database on Ideology, Money in Politics, and Elections (DIME). The measure relies on comparing the mix of campaign contributions that candidates, both winners but also losers, receive from donors. Although some donors, mainly organized interest groups, donate in a strategic manner, the vast majority of contributions are driven by ideological concerns (Bonica 2013, 2014; Snyder 1992). As a result, contribution-based estimates of candidate ideology do a good job of predicting how candidates will behave in the legislature if elected.

The contribution-based scalings, in contrast to roll-call based scalings, allow for the study both winners and losers—not just incumbents who cast roll-call votes. Inevitably, contribution-based estimates are limited to the set of candidates who raise sufficient money to be scaled. Candidates who never raise any money obviously cannot be scaled on the basis of campaign contributions. However, as a practical matter, my resulting focus on “viable” candidates will not limit the study significantly. Any candidate with any chance of winning electoral office receives sufficient contributions to receive a scaling.

The main contribution-based estimate that DIME makes available is the regular CFScore, which pools all of the contributions candidates receive in their entire career, producing the most precise possible estimate of their ideology. The CFScore ranges from roughly -5 to +5 for U.S. House candidates, with smaller

(more negative) numbers indicating liberal positions and larger (more positive) numbers indicating conservative positions. I utilize this measure in cases where I am not studying dynamic parts of the electoral process. In order to investigate possible changes in candidate positions over time, however, I rely on *dynamic* CFScores—the contribution-based estimates of ideology made available in DIME that are estimated separately for each election cycle.

First, I examine how U.S. House candidates position themselves in their first and second election attempts, respectively (regardless of whether they win office or not). Figure 3.1 plots, for each candidate in the dataset who runs in at least two elections, their positions in their first two campaigns. Perfectly horizontal lines connecting each candidates' two datapoints indicate no change in position. As the figure shows, the vast majority of lines are perfectly horizontal. Though there are some changes—reflected in the scattered non-horizontal lines in the plot—they are few and far between.

Figure 3.2 examines this another way. Here, I calculate the change in position between the first and second campaign for every candidate that runs at least twice in the U.S. House. I then plot the distribution of those changes. A huge amount of the mass of the distribution is placed in the immediate vicinity of zero; most candidates simply do not change their positions.

These graphical analyses suggest that candidate positions are stable. Another way to see this more formally is to estimate autoregressive equations of the form

$$\text{Dynamic CFScore}_{it} = \beta_0 + \beta_1 \text{Dynamic CFScore}_{i,t-1} + \epsilon_{it}, \quad (3.1)$$

where for each candidate i we include her previous cycle's CFScore as the regressor. An estimate for β_1 close to 1 suggests a high degree of autocorrelation in candidate positions over time. Figure 3.3 presents the resulting estimate. The black points present raw data, comparing each candidate's position in the election at $t - 1$ (on the horizontal axis) to the candidate's position in the election at t (on the vertical axis). The red line overlaid on the plot presents the regression estimate. As can be seen in the plot, there is an extremely high association between lagged candidate positions and current candidate positions. Again, we see that candidates do not appear to change their positions over time.

Figure 3.1 – The Rigidity of Candidate Positions: Examining Candidates’ First and Second Campaigns. Plots estimates of first- and second-term candidate ideology based on dynamic CFScores for U.S. House Candidates, 1980–2012, who run for election at least twice. As the horizontal lines show, candidates’ estimated positions rarely change.




Figure 3.2 – The Rigidity of Candidate Positions: Examining Candidates’ First and Second Campaigns. Plots estimated change in candidate ideology based on dynamic CF-Scores for U.S. House Candidates, 1980–2012, who run for election at least twice. The vast majority of candidates display almost no change in their ideology.





Figure 3.3 – The Rigidity of Candidate Positions: Evidence from Candidate Careers. Compares dynamic CF-Scores for a candidate in an election at time t to her previous dynamic score at $t - 1$, for U.S. House Candidates, 1980–2012. Previous ideology is an extremely strong predictor of current ideology.



Note: Red line from OLS estimate of equation 3.1.

3.2 Incumbents Do Not Adjust Positions Based On Primary Challenges

Thus far I have presented descriptive evidence that candidate positions do not appear to move much, if at all, over the course of candidates' careers. In this section, I focus more on a causal analysis, and I connect the concept of candidate rigidity to the polarization literature by focusing on one commonly proposed source of polarization in U.S. legislatures: primary challenges. An important body of work discusses the possible consequences of primary challenges to incumbents, especially when members of Congress “get primaried,” i.e., get challenged by more ideologically extreme candidates from their own party (Boatright 2013).¹

One key prediction in this literature is that incumbents should respond to such challenges by moving farther away from the median voter in the district, and closer instead to the median primary voter. Here I test this particular prediction by examining how incumbents respond, ideologically, to the emergence of more extreme primary challengers. Consistent with the descriptive evidence above, I again find that incumbents hold remarkably rigid positions. When challenged by more ideologically extreme primary candidates, incumbents on average make no alterations to their ideological platforms. Though this analysis does not test theories that predict that incumbents *pre-emptively* take non-median views in order to dissuade primary challenges, the results certainly suggest that incumbents' positions—however chosen—do not respond to the ideological positions of opponents.

Obtaining empirical leverage on this question is difficult. First, we must identify “extremist” primary challengers. Then, we must find a way to assess the effects these challengers have on incumbent positions. Simply comparing average incumbent positions when there is or is not an extremist challenger will not accomplish this goal; incumbents who get challenged may differ, systematically, from those who are not challenged, and so, too, may the districts in which the two types serve. If we focus on how individual incumbents *change* their positions when faced with an extremist challenger vs. when faced with a moderate challenger

¹For a review of this literature see Section 2.2.

vs. when faced with no challenger, we can attempt to address this source of bias, because we hold fix both the underlying “type” of the candidate and the district.

However, this “within-incumbent” design still has problems. Most obviously, changes in an incumbent’s platform over time may not be due only to whether or not an extremist challenger is present. We can address this type of dynamic problem by adding year fixed effects to the analysis, performing a difference-in-differences in which we examine the differential way in which incumbents challenged by extremists change their positions relative to incumbents who are not challenged. The difference-in-differences is an improvement over the simple within-incumbent design, addressing both the unobserved heterogeneity across incumbents and districts and across time periods. As a result it offers a plausible look at how incumbents react to extremist challengers.

To implement the difference-in-differences, I estimate equations of the form

$$\begin{aligned} \text{Dynamic CFScore}_{it} = & \beta_0 + \beta_1 \text{Extremist Challenger}_{it} \\ & + \text{Moderate Challenger}_{it} + \gamma_i + \delta_t + \epsilon_{it}, \end{aligned} \quad (3.2)$$

where $\text{Extremist Challenger}_{it}$ is an indicator variable for the presence of an extremist challenger to incumbent i in the primary at time t .

Extremist challengers are identified as follows. For Democratic incumbents, a challenger is coded as “extreme” if her estimated dynamic CFScore is farther to the left (more negative) than the incumbent’s dynamic CFScore in the previous electoral cycle. For Republican incumbents, a challenger is coded as extreme if her estimated dynamic CFScore is farther to the right (more positive) than the incumbent’s dynamic CFScore in the previous electoral cycle. Moderate challengers are constructed analogously, except for candidates farther to the right than the incumbent’s previous position, in Democratic primaries, and farther to the left in Republican primaries.

Table 3.1 presents the estimated results. The first row presents the coefficient estimates on *Extremist Challenger*. In both parties, incumbents exhibit no discernible ideological shift in response to primary challenges from their ideological flanks. Consider the estimate in the first column, for Democrats. Democratic incumbents who enjoy uncontested primaries have an average dynamic CFScore

Table 3.1 – Incumbents Do Not Shift Positions in Response to Extremist Primary Challenges.

	Democrats CFScore	Republicans CFScore	Both Parties Abs CFScore
Extremist Challenger	-0.003 (0.013)	-0.015 (0.014)	0.003 (0.010)
Moderate Challenger	-0.029 (0.013)	0.008 (0.013)	0.020 (0.010)
Intercept (No Challenger)	-0.478 (0.017)	0.969 (0.023)	0.713 (0.015)
N	3,481	2,778	6,259
Year Fixed Effects	Yes	Yes	Yes
Incumbent Fixed Effects	Yes	Yes	Yes

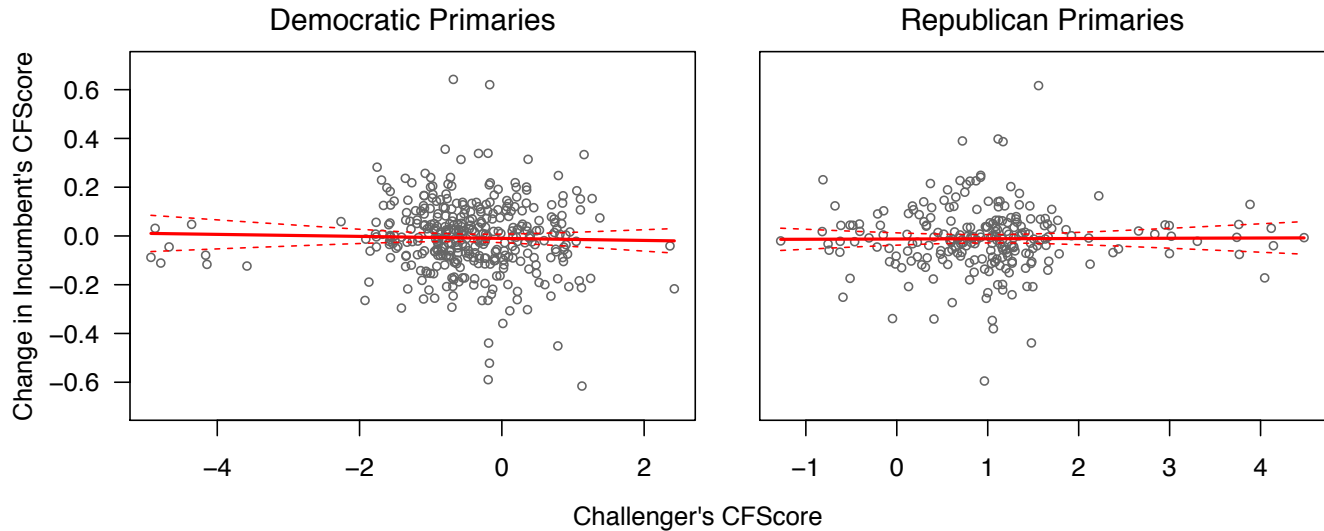
Robust standard errors clustered by incumbent in parentheses. CFSCores are dynamic scores.

of -0.478; those with an extremist challenge have, on average, a dynamic CF-Score 0.003 points less—a minuscule difference, and one we cannot statistically differentiate from the null hypothesis of no difference.

The second column shows similar results for Republican incumbents. Again, the smallest of differences is seen between uncontested incumbents and those facing an extremist challenger. The third row combines the results from the first two columns by taking the absolute value of the dynamic CFScores. Across both parties, incumbents have on average a score of 0.713; those with extremist challengers shift their positions by an estimated 0.003 points, i.e., they do not change at all.

The second row presents the analogous estimates for the case of moderate challengers. Again we find no meaningful changes. Although we can reject the null hypothesis in the first and third columns, the estimated effects are substantively negligible. Consider the point estimate on “Moderate Challenger” in the first column, for example. The intercept in the first column tells us that an uncontested Democratic incumbent has, on average, a dynamic CFScore of -0.478,

Figure 3.4 – Incumbents Do Not Respond to Primary Challenger Positions. The plots present how much incumbents of each party change their ideological positions in response to varying positions of their primary challengers. There appears to be no relationship between challenger position and changes in incumbent positions.



Note: Red lines are OLS fits to data. Dashed lines are 95% confidence intervals. CFscores are dynamic scores.

while those challenged by a more moderate candidate change their position, on average, to a dynamic CFscore of -0.507, a change of -0.029 points. The average distance between the parties' uncontested incumbents, for comparison, is 1.447. This change is thus only 2% the size of the distance between the parties.

Figure 3.4 offers these results in a different way. Each of the two plots compares challengers' dynamic CFscores to the *change* in incumbent dynamic CFscores from the previous election to this one. The overlaid red lines represent simple OLS estimates. For both parties the lines are remarkably flat. Incumbents appear not to change their positions at all regardless of where their challengers lie, ideologically.

The method of identifying extremist or moderate candidates that I have used in these analyses has two potential, related problems. First, some challengers

may have ideological positions very much like those of the incumbent. For these races, we might expect the incumbent to move less, if at all, due to the proximity. Second, since these positions are estimated with noise, challengers close to incumbents will be more likely to be mis-classified as “moderate” challengers when they are in fact extremist challengers, or vice-versa. To address this issue, I follow the technique of Hall (N.d.). In the Appendix to this chapter, I re-estimate the same regressions only using the subset of data in which the distance between the incumbent and her challenger is at or above the median such distance across all races. This ensures that I only perform the analysis on incumbents and challengers with quite distinct positions. I continue to find exactly the same pattern of evidence.

If candidates possessed fluid ideological positions, then they should respond, spatially, to the positions of their opponents. This logic is at the heart of models that predict that incumbents should address primary challenges by moving their positions. If, on the other hand, candidates have rigid positions, then such challenges should have no effect on incumbent positions. The evidence accords with this hypothesis. Incumbents do not respond to moderate or extreme challengers by moving at all. The easiest way to explain this behavior, in conjunction with the previous sections’ results on candidate positions over time, is that candidates possess rigid, rather than fluid, positions.

3.3 Candidate Identity Linked to Candidate Positions

If, as I argue, candidates come pre-committed to certain ideological positions, then we might expect to observe systematic differences in candidate positions across salient aspects of candidate identity. The purely Downsian world predicts all candidates—regardless of backgrounds, identities, etc.—to move to the middle. Thus we should observe, under this theory, no systematic differences candidate positions across candidate identities, *for a set of candidates running to represent the same voters*. This last caveat is crucial. We might observe systematic differences in, for example, men and women candidates across the pooled

dataset. But these differences could indicate that men and women candidates arise in different areas, where voter preferences are different, and thus would not necessarily run counter to Downsian predictions. But if candidates offering to represent the *same* median voter offer systematically different positions based on observable demographic characteristics, then we might conclude there is more to candidates' positions than the location of their districts' median voters.

To examine this hypothesis, I focus on a single demographic trait which is readily observable in the data: candidate gender. An existing literature studies differences in ideology by gender, both for candidates and for elected legislators (see for example McDermott 1997; Welch 1985); I simply reinforce the findings in this literature using a larger dataset containing the full universe of U.S. House candidates. I use the Bonica (2014) dataset, which includes an estimate of each candidate's gender based on first names. First, in Figure 3.5, I plot the distribution of candidate ideology for men and women by party. By splitting by party I ensure that observed differences between men and women candidates do not simply reflect an imbalance in their prevalence across parties—e.g., since women are more frequently Democratic candidates (Thomsen N.d.), finding that women on average have more liberal positions than men might only reflect differences across the parties and not across men and women candidates in other ways. As the densities show, women offer on average more liberal positions. This difference is especially marked in the Democratic party but present in both parties.

It is still possible that the observed difference reflects a difference in context; within party, women may run for office in different types of districts than do men, which would again prevent us from concluding that we are detecting a demographic difference. Accordingly, I estimate regression equations of the form

$$CFScore_{ijt} = \beta_0 + \beta_1 Woman_{ijt} + \beta_2 Dem_{ijt} + X_{jt} + \epsilon_{ijt}, \quad (3.3)$$

where $Woman_{ijt}$ is an indicator variable for candidate i in district j in the election at time t being a woman and Dem is the corresponding indicator variable for party. The variable X_{ijt} stands in for a vector of district-level control variables to address the possibility of differences across the set of districts that see women candidates more often.

Figure 3.5 – Candidate Ideology by Gender and Party.
 Among candidates, women hold, on average, more liberal positions than men of the same party.

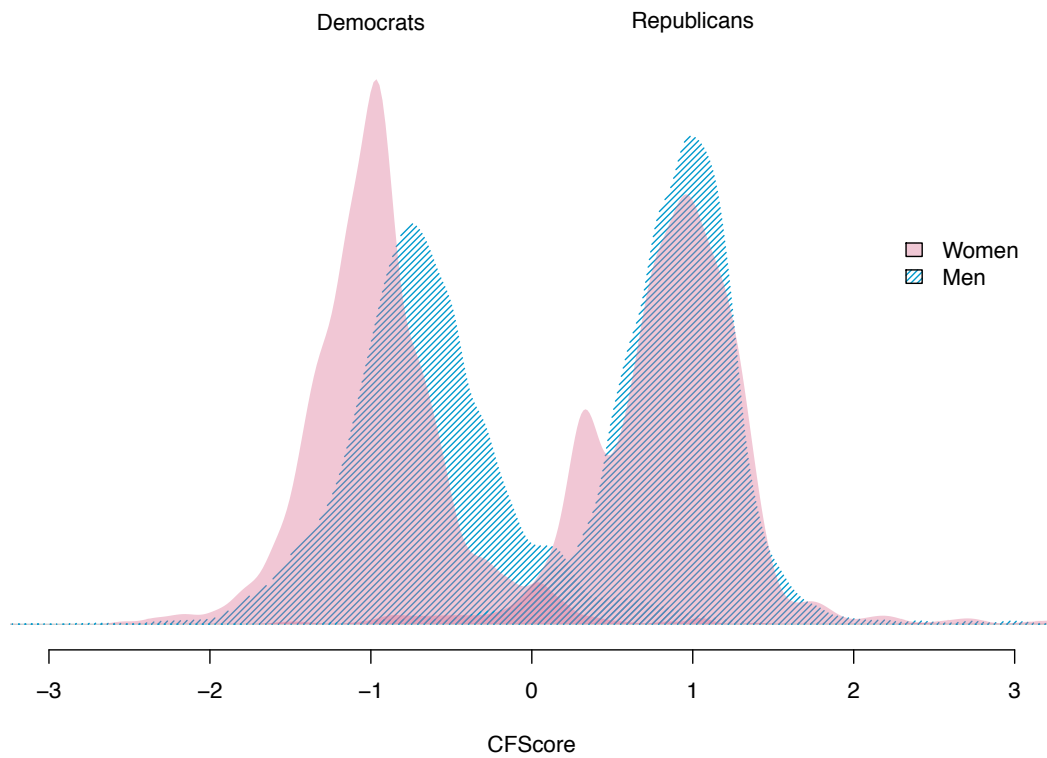


Table 3.2 presents the results using four different specifications. In the first column, I include no additional controls—thus computing the simple average difference in ideology across men and women candidates within party. In the second column, I add district-level controls for the average donor ideology and the district’s previous Democratic vote share for president. In the third and fourth columns I repeat this exercise but with the addition of district and year fixed effects. Across all four columns we see a steady difference; women candidates are consistently estimated to be, on average, more liberal than male candidates.

There are many possible reasons for this relationship, and understanding it is far beyond the scope of this study. But the documented pattern has important implications for this book’s argument. Candidates are not just strategic, blank-

Table 3.2 – Women Candidates Hold More Liberal Positions Than Male Candidates, On Average.

	CFScore	CFScore	CFScore	CFScore
Woman	-0.20 (0.02)	-0.17 (0.02)	-0.16 (0.02)	-0.16 (0.02)
Democrat	-1.62 (0.02)	-1.75 (0.02)	-1.65 (0.02)	-1.76 (0.02)
N	21,178	14,545	21,178	14,545
Year Fixed Effects	No	No	Yes	Yes
District Fixed Effects	No	No	Yes	Yes
Controls	No	Yes	No	Yes

Robust standard errors clustered by district in parentheses. Controls in columns 2 and 4 are average donor CFScore by district and district’s previous Democratic presidential vote share.

slate agents seeking out the median voter in their district. They come to elections with pre-existing views that arise, in part, from pre-existing traits, like identities and experiences. They do not change these views on a whim. As a result, who is willing to run for office is a crucial factor in determining the positions voters have the opportunity to choose between.

3.4 When (If Ever) Do Candidates Change Positions?

Candidate positions are rigid, on the whole. But that does not mean they have no room to maneuver. The results presented here have all been “on average” findings, averaging over candidates, eras and, most importantly, issues. In the big picture, candidates rarely change their positions, but we should not conclude that candidates *never* change positions on *any* issues. To do so would be to caricature a much more nuanced process.

For example, several recent papers explore conditions under which incumbents change their positions. Feigenbaum and Hall (2015) shows how localized, negative shocks from trade cause U.S. House incumbents to vote more protectionist on trade bills while leaving their positions on other issue areas unchanged—especially when they are more threatened, electorally. Clinton and Enamorado (2014) documents how MCs change their roll-call records in response to the roll-out of Fox News in their districts. Both papers uncover unusually salient situations in which incumbents will change their positions, but because effect sizes are relatively modest, it is unsurprising that such changes do not change the larger picture of rigid positions much.

Another way in which incumbents can alter their overall platform is by taking action on specific issues. Sulkin (2005) argues that incumbents “take up” issues that their challengers prioritize in unsuccessful bids to unseat them. This represents an important dimension of flexibility for incumbents, but it is distinct from changing ideological positions. It focuses on specific policy actions rather than incumbents’ overall ideological portfolio, and it depends on action rather than position.

Candidates have many reasons to alter their positions, both for strategic gain—if they can do so while escaping punishment for “flip-flopping”—and due to changing personal preferences. The candidate-supply theory does not rule out such changes, but it does argue that they should be rare. The more common these changes are, the less binding a constraint the candidate supply is likely to be. The more malleable candidates are, the more voters can pressure them to cater to their demands, the less it matters, from an ideological perspective, who runs for office. But, as the empirical analyses in this chapter have shown, such changes are in fact exceedingly rare. Candidates positions appear to be quite rigid, and voters are thus left to select candidates with pre-existing positions rather than to pressure candidates into catering to their preferred ones.

Summary

A vast literature in American politics is based on a theoretical view of fluid candidate positions. How fluid are candidate positions, empirically? In this

chapter I have offered a series of analyses that suggest that candidate positions are highly rigid, not fluid. Within their careers, candidates rarely, if ever, display any marked shifts in their ideology. This rigidity persists even in the face of primary challenges from incumbents' ideological flanks—often thought to contribute to polarization. Incumbents display no change in their positions, on average, in response to challenges by extremist primary challengers (or, for that matter, in response to moderate primary challengers).

The inflexibility of candidates' positions is consistent with a candidate-supply theory in which candidates have underlying traits, including their ideology, that spring from longstanding, personal characteristics and experiences. Consistent with this idea, the latter part of this chapter echoes previous literature in showing how candidate ideology varies, systematically, with an important component of candidate identity, sex. Women candidates in both parties offer systematically more liberal positions. This is true even when making comparisons among candidates running in the same district, thus highlighting the role of identity, rather than purely strategic concerns, in determining the positions that candidates offer to voters.

By showing that candidate positions are rigid, I have established that the candidate supply *could* matter. If the types of candidates that voters would most prefer to support do not run for office, then the candidate supply becomes a *binding constraint* on voters. On the other hand, if the types of candidates that voters most prefer do run, then the candidate supply does not bind. Now, in Chapter 4, I turn to evidence that the candidate supply does, in fact, prevent voters from electing more moderate candidates.

Appendix to Chapter 3

Table 3.3 – Incumbents Do Not Shift Positions in Response to Extremist Primary Challenges: Including Only Ideologically Distinct Challengers.

	Democrats CFScore	Republicans CFScore	Both Parties Abs CFScore
Extremist Challenger	0.022 (0.017)	-0.044 (0.017)	-0.020 (0.014)
Moderate Challenger	-0.064 (0.016)	0.027 (0.019)	0.039 (0.014)
Intercept (No Challenger)	-0.477 (0.018)	0.965 (0.024)	0.713 (0.015)
N	3,280	2,644	5,924
Year Fixed Effects	Yes	Yes	Yes
Incumbent Fixed Effects	Yes	Yes	Yes

Robust standard errors clustered by incumbent in parentheses. CFSCores are dynamic scores.

Chapter 4

The Candidate Supply and Constrained Voters

So far we have only explored the nature of the candidate supply, without regard to voter demand. Rigid positions mean that voters must choose candidates for office carefully; the positions candidates offer today are likely to represent the positions they will offer down the line, too. As a result, elections *select* candidates based on ideology, potentially, but they do not *pressure* incumbents into changing positions. In this chapter, I therefore examine the process by which elections select candidates for office. Once we measure voter demand for candidates, we can compare it to candidate supply to see whether or not the supply appears to constrain the choices of voters in a meaningful way.

4.1 Voters Choose More Moderate Candidates When They Can

Note: The contents of this section are based on Hall (N.d.).

When given the choice, do voters actually prefer more moderate candidates? Several empirical papers suggest that, in the contest of the U.S. House, the answer is yes. Canes-Wrone, Brady, and Cogan (2002) shows that incumbents with more moderate roll-call voting records perform better, on average, than incumbents with less moderate records. Erikson and Wright (2000), Ansolabehere, Snyder, and Stewart (2001), and Burden (2004) present similar results. I extend these results in two ways. First, I incorporate all candidates, not just incumbents. Second, I employ a regression discontinuity design to ensure that the underlying districts in which more extreme or moderate candidates run are otherwise identical. Specifically, I examine the effect the “as-if” random nomination of extremist

vs. moderate candidates in close primary elections on electoral outcomes in the general election. Consistent with the previous literature, I again find that voters strongly prefer more moderate candidates, on average.

For information on classifying candidates as extremists, I deviate from the approach of the previous chapter. Here, I rely on the estimated ideological positions of U.S. House primary candidates from Hall and Snyder (2014), which covers the years 1980–2010.¹ Technical details on the method are available in Appendix B. Candidates are scaled on the basis of their primary-election campaign receipts,² imputing each candidate’s ideological position from the contribution-weighted average estimated positions of her donors.³ The donors’ positions are estimated as the contribution-weighted average DW-NOMINATE score (Poole and Rosenthal 1985) of incumbents they have donated to, but excluding donations to candidate i when computing the score for each candidate i . While very similar to CFScores, I choose this technique for this analysis only because it produces scalings that are specific to the primary election. Using general-election contributions—though statistically efficient for many other purposes, like those in Chapter 3—would introduce post-treatment bias in the present setting since the “treatment,” the nomination of the extremist, occurs prior to the start of the general-election campaign.

To validate this scaling technique in the sample used for analysis, Figure 4.1 compares it to observed DW-NOMINATE scores for primary candidates who go on to win the general election. The donor-based scaling of candidates correlates with observed DW-NOMINATE scores at 0.90.⁴ This is consistent with a fuller battery of validation tests presented in Hall and Snyder (2014). As a result,


¹I use the scalings that employ a cutoff of 10 on the number of unique contributions a donor must make in order to enter the scaling procedure.

²Hall and Snyder (2014) removes all contributions classified for the general election, according to the FEC’s disclosure requirements, and also removes contributions received after the date of the primary election.

³In practice, weighting these donations by dollar size makes little difference in the scalings.

⁴The within-party correlations are 0.66 for Democrats and 0.56 for Republicans. These within-party correlations are very similar to those in Bonica (2013), which reports within-party correlations with DW-NOMINATE of 0.66 for Democrats and 0.64 for Republicans analyzing only sitting incumbents. The within correlation for Republicans is slightly lower in the current sample, probably because the comparison includes some challengers who go on to be incumbents later. Bonica (2013) reports a within-party correlation of 0.49 for these Republican candidates.

Figure 4.1 – Estimated Ideology of Primary Candidates and Observed Roll-Call Behavior. For primary candidates who go on to win the general election, the contribution-based estimate of their ideology accords well with how they vote on roll calls.



Note: Points labeled ‘D’ represent Democratic incumbents; those labeled ‘R’ are for Republican incumbents.

there is good reason to believe that the estimated primary candidate positions are reflective of their actual ideological positioning. Importantly, any random error in these estimated positions biases the subsequent analysis against finding differences in outcomes for more and less extreme candidates.

In primary races with two major candidates, the race is tentatively identified as being between an extremist and a relatively moderate candidate if the difference between their estimated ideological positions is at or above the median in the distribution of ideological distances between the top two candidates in all contested primary elections. This median distance translates to roughly one third the distance on the DW-NOMINATE scale between the medians of the two parties in the 112th Congress, and it is approximately two to three times as large

as the average distance between representatives and their own party’s median.⁵ These are therefore races between candidates who offer meaningfully different platforms within the umbrella of their party.

Using a strong cutoff like this has two potential advantages, as I discussed in a previous analysis in Chapter 3. First, it may reduce the number of incorrect moderate/extreme labels caused by measurement error in the donor scores. Second, it ensures that we are focusing on strong comparisons in which the two primary candidates are starkly different.

Data on U.S. House primary and general elections is compiled from primary sources by Ansolabehere et al. (2010). I focus on elections in the years 1980–2010 to match the data on candidate positions. I keep all primary elections in which at least two candidates have donor scores.⁶ Among these elections, I analyze the two candidates with the top two vote totals, and I calculate each candidate’s share of the top two vote total.⁷

I estimate models of the form

$$Y_{it} = \beta_0 + \beta_1 \textit{Extremist Primary Win}_{it} + f(V_{it}) + \epsilon_{it}, \quad (4.1)$$

where *Extremist Primary Win*_{it} is an indicator variable for the extremist winning the primary in district *i* at time *t*. Thus β_1 is the quantity of interest, the RD estimator for causal effects from the “as-if” random assignment of an extremist in the general election.⁸ The variable Y_{it} stands in for three main outcome variables:

⁵These calculations are performed as follows. First, I take the median primary-election distance between candidates, which is roughly 0.109. To convert this number to the DW-NOMINATE scale, I plug it in as *X* in a regression predicting DW-NOMINATE based on the contribution-based scaling. I then compare this converted number (0.32) to the observed difference in party medians in the 112th Congress (1.072), and the observed mean differences between each representative and his or her party’s median (0.14 for the Republicans and 0.1 for the Democrats).

⁶There are 504 such elections in the dataset.

⁷For the few cases of runoff primaries, I include the candidates and vote shares of the two-candidate runoff election.

⁸The use of the RD to obtain exogenous variation in candidate types rather than party (Democrat vs. Republican) is similar in spirit to several previous studies. Brookman (2009) uses such a design to study the effects of copartisan incumbency in Congress on presidential elections returns. Broлло and Troiano (2012), Brookman (2014*b*), and Ferreira and Gyourko (2014) examine the effects of the “as-if” random assignment of female vs. male candidates to office.

party vote share, party victory, and the DW-NOMINATE score of the winning general-election candidate in the ensuing Congress. The term $f(V_{it})$ represents a flexible function of the running variable, the extremist candidate's vote-share winning margin, i.e., the extremist candidate's share of the top two candidates' vote less 0.5, which determines treatment status. I present estimates using a variety of specifications for f as well as at different bandwidths, following the usual RD practices (Imbens and Lemieux 2008). Typically, f either contains a high-order polynomial of the running variable or a local linear specification estimated separately on each side of the discontinuity.


The key identifying assumption of the RD is that potential outcomes are smooth across the discontinuity, i.e., that districts where the relatively moderate primary candidate barely wins (or, equivalently, those where the extremist candidate barely loses) are in the limit comparable to those in which the relative moderate barely loses (or, equivalently, those where the extremist candidate barely wins). Assumptions of “no sorting” like this have been challenged in the context of general-election U.S. House races (Caughey and Sekhon 2011; Grimmer et al. 2012; Snyder 2005), but across numerous other electoral contexts are found to be highly plausible (Eggers et al. 2015).

For primary elections, this assumption is extremely plausible; to sort, candidates would need to have precise information about the expected outcome of primary elections, and would need to exert extra effort only after finding out that the election was going to be extremely close.⁹ Given the difficulty even general-election campaigns have in predicting votes (Enos and Hersh 2013), and the typical “running scared” mentality that leads candidates to pull out all the stops (King 1997), this seems unlikely. In addition, I validate the assumption in the Appendix by presenting balance tests using the same samples and specifications as the main results in the paper. No evidence of sorting is found.¹⁰

⁹Alternatively, barely winning extremists might have some ability to alter vote totals after the election in a way that barely losing extremists cannot. This seems even less likely.

¹⁰Of 24 balance tests run, only 2 (8%) reject the null of no difference. These rejections are for one specification of the W-NOMINATE score and one specification of lagged electoral victory, but do not persist across alternate specifications for these variables. See the Appendix for more details.

Figure 4.2 – General-Election Vote Share After Close Primary Elections Between Moderates and Extremists: U.S. House, 1980–2010. The close election of the more extreme primary candidate causes a decrease in general-election vote share for the party.



Note: Large black points are averages in 0.02 point bins of the relatively extreme candidate’s winning margin; small gray points are raw data. Lines are OLS fits from raw data estimated separately on each side of threshold. Average general-election vote shares are above 0.5 on both sides of the discontinuity because contested primaries are more likely to occur in districts where the normal vote is tilted towards the party.

In addition to testing for sorting, it is also important to show that the RD estimate is not sensitive to choices over the size of the bandwidth and the functional form of the running variable. In the Appendix, I replicate the main analysis on general-election outcomes at a large variety of bandwidths and specifications and show that the resulting estimates consistently produce the same conclusion.

Figure 4.2 plots the discontinuity in the data. As can be seen, when the extremist goes from barely losing the primary to barely winning it (horizontal axis), the party’s general-election vote share decreases noticeably. Among “coin-flip” primary elections between a relative moderate and an extremist, the nomination

Table 4.1 – RDD Estimates of the Effect of Nominating an Extreme Candidate on General Election Vote Share, U.S. House 1980–2010.

	Vote Share	Vote Share	Vote Share	Victory	Victory	Victory
Extremist Win	-0.12 (0.06)	-0.10 (0.03)	-0.08 [0.04]	-0.53 (0.22)	-0.37 (0.11)	-0.35 [0.17]
N	83	252	135	83	252	148
RDD Bandwidth	5	–	8.51	5	–	9.68
Specification	Local Linear	Cubic	IK	Local Linear	Cubic	IK

Maximum of robust and conventional standard errors in parentheses. Columns 3 and 6 use optimal bandwidth technique from Imbens and Kalyanaraman, implemented using `rdob` in Stata. Standard errors from this procedure in brackets.


of the extremist appears to cause a large decrease in the party’s general-election vote share.

Table 4.1 presents the estimates from equation 4.1 using general-election vote share and victory as the outcome variables, with three specifications for each. In the first and fourth columns, I use a 5% bandwidth and a local-linear specification of the forcing variable, estimated separately on each side of the discontinuity. In the second and fifth columns, I use all the data and include a cubic specification of the running variable. In the third and sixth columns, I employ the “optimal bandwidth” procedure from Imbens and Kalyanaraman (2012).¹¹

As the table shows, the “as-if” random assignment of the extremist to the general causes approximately a 9–13 percentage-point decrease in the party’s share of the general-election vote, and a 35–54 percentage-point decrease in its probability of victory. These are large effects. This analysis has thus established the electoral penalty that the party faces when it is randomly assigned an extremist candidate for the general election. General-election voters prefer more moderate candidates by a large margin.

¹¹Because sample sizes are relatively small, especially at small bandwidths, I report the maximum of robust and conventional standard errors for the local linear and cubic specifications (Angrist and Pischke 2009). I report the standard errors from Imbens and Kalyanaraman (2012), as reported by the procedure, for the optimal bandwidth specifications.

Figure 4.3 – Comparing RD Sample to Full Population of Congressional Districts. The districts in the RD sample used to estimate the penalty to extremists are extremely similar to the full population of districts in terms of partisanship, as measured by presidential vote share.



These results depend on a particular subset of all races—those with close primary elections between a more moderate and a more extreme candidate. The argument I want to make is about the preferences of voters more generally. We need to consider, therefore, whether or not these results are applicable to a broader swath of congressional districts before we move on.

Since the penalty is much smaller in safe districts, if safe districts are a larger proportion of all districts than of those in the RD sample, then the RD results will not generalize to all districts. However, this is not the case. Figure 4.3 compares the normal vote for districts in the RD sample to all congressional districts over the same time period; as the plot shows, the RD sample is actually extremely similar, in terms of partisanship, to the full population of districts. The penalty

to extremists I document in the primary RD is therefore likely to be a more general phenomenon in congressional elections.

4.2 The Candidate Supply Constrains Voters

When given the opportunity, voters tend to support more moderate candidates. Yet, at the same time, we know that incumbents—those candidates elected to office—have been offering increasingly extreme positions to voters over time. How do we square these two facts? Voters may prefer moderate candidates, but they can only elect such candidates if they run for office. The candidate-supply theory posits that the explanation is that moderate candidates are simply not running for office. In this section, I verify this key empirical component of the argument.


To do so, I examine the ideology of the entire supply of candidates, again using the CFScore scalings based on the mix of campaign contributions candidates received.¹² Specifically, for the 22,727 Democratic and Republican U.S. House candidates in the dataset, I first regress their CFScores on a set of dummies for each congressional district, in order to account for differences in positions across districts. I then take the average of the residualized candidate positions for each year by party. I plot these averages in Figure 4.4.

The figure reveals two important patterns about the candidate supply. First, as the points at the far left show, for the year 1980, the two parties' candidate supplies have always diverged significantly, ideologically. Second, the degree to which the two parties candidate supplies diverge has grown, markedly, over time. Indeed, the average distance between the candidates of the two parties has grown from 1.2 points on the CFScore scale, in 1980, to just over 2.0 points in 2010. The ideological divergence between the two parties' candidate supplies has thus almost doubled in a 30-year period.

Both of these patterns are consistent with the measures of legislative polarization—based on incumbent roll-call voting—which also reveal a pronounced and growing divergence over this time period (see McCarty, Poole, and Rosenthal 2006). Figure 4.5 illustrates this dynamic. The plot shows the absolute distance between

¹²For a discussion of this measure and its validity, refer back to Chapter 3.

Figure 4.4 – Average Ideology of All U.S. House Candidates by Party, 1980-2010. In both parties, those running for the U.S. House have become more ideologically extreme over time.



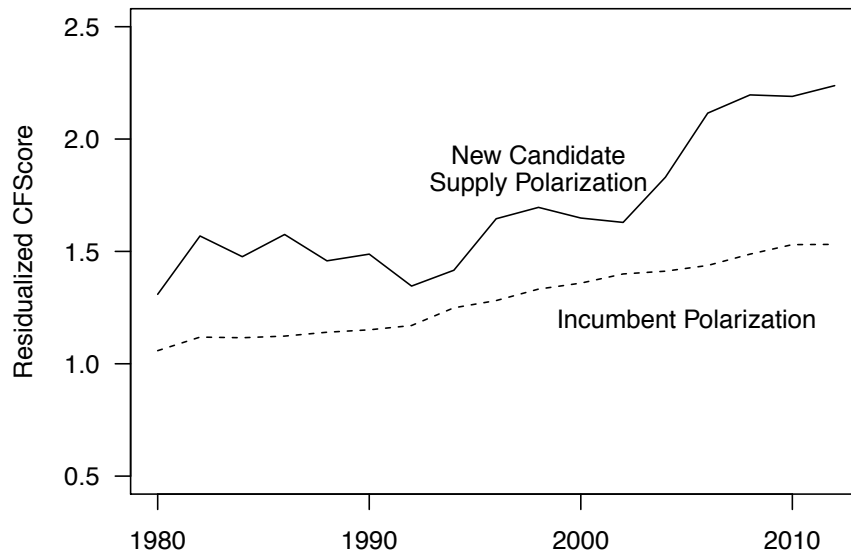
Note: Points represent yearly averages by party; lines are OLS fits to raw data. Candidate ideology measured by static CFScores residualized by congressional district. Standard error lines omitted because they are too small to discern.

the average ideological position of each party for each year, separately for two groups of candidates: incumbents and challengers/open-seat candidates, whom I call “new” candidates. Like before, candidate ideologies are first residualized by district so that the resulting calculations reveal candidate divergence and not sorting across districts.

The two lines track each other well ($r = 0.9$). As incumbent polarization has risen so, too, has polarization in the new candidate supply. Incumbent polarization is also consistently below that of new candidate polarization. The results in section 4.1 provide the likely explanation for this fact. Voters select for more moderate candidates from among the pool—thus producing Democratic and Republican candidates who are closer together, on average, than the overall candidate pools. Nonetheless, incumbent polarization remains quite high. The

Figure 4.5 – Average Polarization of U.S. House Incumbents and New Candidates, Respectively, 1980-2010.

Each line is the absolute difference in average CFScores across the two parties for the relevant candidate group, by year. Incumbent polarization has grown at the same time as has the candidate supply.




Note: Candidate ideology measured by static CFScores residualized by congressional district.

candidate supply does not give voters the opportunity to shrink the difference between the candidates of the two parties further.

Correlating two time series is hardly a way to establish a causal relationship, but the fact that candidate positions are rigid, as Chapter 3 established, permits us to interpret these relationships to some degree. Voters can only elect those who run for office, and those who run cannot, according to the candidate-supply theory, adjust their positions to respond to voter demand. Thus, if the candidate supply polarizes, the ideological composition of the legislature is forced to polarize as a result.

The descriptive analyses above reflect average ideologies. Although they use candidate ideologies that are residualized by district, there is still the possibility that *average* candidate ideologies are changing mostly due to extreme outliers.

Figure 4.6 – Distribution of the District-level Ideological Difference Between the Most Moderate Republican and Democratic Candidates Across Eras. While already a polarized era, more districts had choices between ideologically similar candidates in the 1980s than they have in the 2000s.



Even though there may be *fewer* moderates, and even though the candidate supply might be on average more ideologically extreme, it is still possible that a sufficient number of moderates remain such that the candidate supply is not so strong a constraint as I suggested.

To address this possibility, I now examine differences between Democratic and Republican candidates by district. Specifically, for each district, I calculate the distance between the *right-most* Democrat, i.e., the most moderate Democrat, and the *left-most* Republican, i.e., the most moderate Republican, who enters the race. This includes primary-election winners and losers, since I am investigating the entire candidate supply.

Figure 4.6 plots the resulting distribution of district-level distances between the most moderate Republican and Democrat, for all races in which at least one

Democrat and at least one Republican entered their respective primaries and raised enough money to obtain a CFScore. The gray density represents the distribution of these distances in the 1980s. Although the 1980s were already an era of polarization—as we see reflected by the fact that the majority of districts see significant differences in ideology between their most moderate Republican and most moderate Democratic candidates—there are still a small group of districts who have the option of similarly located Democratic and Republican candidates, as reflected by the positive density towards the left of the plot.

The purple distribution overlaid on the 1980s density depicts the same data for the 2000s. Here, almost no districts ever have the chance to choose between similarly located Republican and Democratic candidates. Very little density is located anywhere near the lefthand side of the plot. Instead, the entire density has shifted markedly to the right, indicating higher average distances between even the most moderate Republican and Democrat in each race. The plot thus reflects the degree to which the candidate supply constrains voters. Opportunities to choose moderate candidates from either party are few and far between.

Summary

In this chapter, I have shown why the candidate supply is a binding constraint on the electoral process. Voters exhibit a consistent preference for candidates who offer more moderate positions, but are faced with a supply of candidates that contains few—and fewer over time—moderates. Because these candidates cannot change their positions, the absence of moderate candidates creates an absence of moderate incumbents. The candidate supply is therefore an important input into the democratic process and an important driver of legislative polarization.

Understanding *why* the candidate supply has shifted in this manner is a primary goal of the candidate-supply theory. Now, in Chapter 5, I show how the costs and benefits of running for office affect the ideological composition of the candidate supply as the theory predicts.