Puppetmasters or Pawns?
The Power of Congressional Staffers

PRELIMINARY WORK. DO NOT DISTRIBUTE.

Andrew Kao
Sara Ji*

March 20, 2025

Abstract

Members of Congress do not write their own bills—their staffers do. We estimate the importance of personal staffers in the US House of Representatives. To do so, we extend the mover design to team settings where an individual's output is unobserved. We find that staffers account for 40% or more of the gap in bills written between high and low-performing offices. Conventional estimation methods are either infeasible, or misspecified and deliver estimates less than half as large. We also find that staffers are more productive when in legislative jobs (except when they are promoted to Chief of Staff), they contribute more in diverse offices, and they are more effective when paired with effective Representatives. Finally, we estimate staffer ideology and find that staffers may moderate ideological extremism within parties and are most impactful on less partisan issues, while Representatives drive the majority of polarization.

Keywords: Congress, staffers, legislation, teams, polarization
JEL Classification: D72, D73, J44

*Kao: Harvard University. Ji: Harvard University. Email: andrewkao@g.harvard.edu. Many appreciated suggestions, critiques and encouragement were provided by Julio Solís Arce, Edward Glaeser, Vincent Pons, Jesse Shapiro, Christopher Walker, Yu-Jia Wang, David Yang, and Hye Young You. We also thank Benjamin Jones for sharing code to replicate Ahmadpoor and Jones (2019). Kao acknowledges financial support from the NSF Graduate Research Fellowship Program and the Harvard Center for American Political Science.

1 Introduction

Congress produces over 10,000 bills a year, passes over 500 of these into law, and controls a 5 trillion dollar budget. Although the media typically highlights Representatives and Senators in the production and passage of bills, their staffers are the ones who draft the actual bills (Fox and Hammond, 1977). However, the role of staffers in the legislative process remains a black box. Some scholars have argued that staffers are no more than "extensions" of their member of Congress (Kingdon, 1989; DeGregorio, 1988), who are "bound" to them (Salisbury and Shepsle, 1981; Hall and Deardorff, 2006), and whose value after leaving Congress is the connection to the specific member they worked for (Blanes i Vidal et al., 2012; Bertrand et al., 2014)—a theory of staffers as "pawns." Others have argued that staffers can act autonomously (Romzek and Utter, 1997), influence legislative productivity (Ommundsen, 2023), and shape policy and originate legislation (Hagedorn, 2015; Crosson et al., 2020)—a theory of staffers as "puppetmasters." Over the past decades, the number of Congressional staffers has often been cut in the name of government efficiency, tied to concerns that staffers focus more on "political strategy over deeper policy discussions."[1] How important are Congressional staffers, and how have they contributed to the production and ideology of legislation? In this paper, we quantitatively estimate the legislative productivity and ideology of personal staffers in the US House of Representatives.

Although some Congressional staffers eventually achieve a degree of fame,[2] the majority of staffers work away from the limelight. Their names do not appear on the bills that they write. This makes it difficult to know whether offices are productive because of their staffers or their Representatives. We make progress on this question by extending the canonical mover design (or AKM model, Abowd et al. (1999)) to team settings where the contribution of each individual is unobserved. Heuristically, we may infer that a staffer is very productive if we observe that office B becomes much more productive and office A much less so in the period after the staffer moves from A to B. With this team-based mover model, we are able to estimate the share of legislative productivity differences between offices that are driven by staffers. Specifically, we estimate the share of differences in bills produced between above (high-performing) and below median (low-performing) offices that is attributable to the team of staffers. We focus our analysis on the introduction of bills, as they are the foundational building blocks of legislation that

[1] `https://www.brookings.edu/articles/vital-stats-congress-has-a-staffing-problem-too/` and `https://washingtonmonthly.com/2014/06/09/the-big-lobotomy/`

[2] For instance, President Lyndon B. Johnson was a Congressional staffer from 1931 to 1937.

staffers are mostly likely to directly contribute to, but we also examine lower frequency outcomes such as the passage of laws or the introduction of important bills. We develop two novel strategies to estimate the model dynamically, allowing us to estimate effects at different time horizons and test for pre-trends. We are also interested in estimating richer, nonlinear models of bill production and staffer ideology. Though this is not possible with OLS, as measurement error in fixed effects do not average out and thus bias estimates, we use Bayesian methods that address these issues and allow us to estimate richer models. This estimation framework can be applied to any setting featuring team production where individual output is unobserved.

We examine the impact of personal staffers to Representatives in the House. Our baseline finding is that these staffers account for 40% of the differences in bills written between above and below median productivity offices. Using a novel dynamic estimation strategy, we find that staffers have an immediate impact upon joining an office that persists until they leave, with no evidence of pre-trends. In an alternative exercise, we dynamically predict when staffers move using nothing but their out-of-sample fixed effect; this validates our individual effect estimates. We validate our baseline findings with a Bayesian approach that estimates the distribution of staffer and Representative effects. These results collectively demonstrate the importance of staffers to the legislative process.

We next investigate sources of staffer productivity using the Bayesian approach, which makes it possible to examine nonlinear interactions and heterogeneity. We first test the extent to which different jobs facilitate the production of bills. Unsurprisingly, we find that staffers in legislative roles, such as Legislative Assistants or Directors, contribute significantly to the production of bills. More surprisingly, we find that when promoted to Chief of Staff, staffers become less productive—so much so that they would contribute more to *legislative* output in *non-legislative*-facing jobs (such as communications). We also find that diversity—across ethnicity, gender, educational background, and prior work experience—offers modest but statistically significant gains in staffer productivity. Next, we estimate a nested CES production function to uncover the relationship between staffers and Representatives. We find that while staffers within an office are highly substitutable, Representatives and their team of staffers have a production function that is much closer to Cobb-Douglas. Staffer productivities are correlated with their wages. This motivates a simple test of potentially Pareto improving trades of staffers between offices. We find that there are very few such Pareto-improving trades in the data, implying that offices hire staffers in a sophisticated manner that accounts for both team and Representative skills.[3] We demonstrate that our Bayesian approach has advantages over the prior liter-

[3]To the best of our knowledge, this test for Pareto improvements in AKM-style models is novel.

ature in estimating this nested CES model. We also investigate misspecified versions of the team-based AKM model that do not fully account for the effect of teams, and show that they deliver estimates of the staffer share only half as large as ours.

Finally, we explore how staffers influence the ideology of the bills they help craft. Along the left-right ideological spectrum (captured by the first DW-Nominate dimension), we find that Representatives account for the majority of the variation, driving polarization across parties. Although staffers contribute slightly to cross-party polarization, they moderate extreme offices within a party, reducing within-party polarization. We show that Representative ideologies are highly correlated with both Trump voteshares and the safeness of their seat, while staffer ideologies are not. Intuitively, because Representatives must face election, their ideology is more closely tied to their constituents', making staffers a potentially moderating force. On less partisan policy issues (captured by the second DW-Nominate dimension), staffers account for a much more substantial 80% of the share in ideological differences.

This paper directly contributes to three strands of literature. First, on the effect that Congressional staffers have on the legislative process. In economics, attention has primarily focused on the "revolving door" phenomenon and the impact of staffers-turned-lobbyists (Blanes i Vidal et al., 2012; Bertrand et al., 2014).[4] Political scientists have long debated the role that staffers play in Congress. Some have claimed that they are little more than extensions of their member of Congress,[5] while others have argued for their ability to independently influence the legislative process.[6] This work has been largely qualitative. Recent studies have focused on correlationally documenting the role of staffers: Leal and Hess (2004) finds that House freshmen are more likely to hire experienced staffers, while Crosson et al. (2020) finds that experienced staff can help new members of Congress and committee chairs, but not the average member of Congress. Ommundsen (2023) finds that legislative output is higher in committees with experienced staff. Our primary contribution is to estimate the individual effect of staffers, and to decompose the differences in legislative output between Congressional offices into the effect of staffers against the effect of offices themselves. Closest to us in this regard is Montgomery and Nyhan (2017), who assume staffers are exogenously assigned conditional on observables and find that Congress members that exchange important staffers are more similar in legislative effectiveness and ideology. In addition to providing a design that does not assume conditionally exogenous assignment at a point in time, we are also able to speak to mechanisms

[4]Other studies on the revolving door include Shepherd and You (2020) and Lazarus et al. (2019).

[5]See, among others, Kingdon (1989); DeGregorio (1988); Salisbury and Shepsle (1981); Hall and Deardorff (2006)

[6]See, among others, Fox and Hammond (1977); Romzek and Utter (1997); Hagedorn (2015).

that make certain staffers more effective than others.

Second, on the literature estimating individual effects in team environments. We extend on a large literature that uses the mover design in value-added models (Abowd et al., 1999),[7] to settings where production occurs in teams and individual contributions are unobserved. We know of four papers that tackle this question: Chan (2021), Bonhomme (2021), Bergeron et al. (2022), and Ahmadpoor and Jones (2019). None estimate a nonlinear model featuring an average team size greater than 2. Other studies, particularly in healthcare, have estimated individual effects in team environments without estimating the effect of an entire team: in Chen (2021), the analysis focuses only on estimating the effects of the 'main' physician on a team, while in Silver (2021), physicians work individually on cases but may be influenced by other nearby physicians via peer effects.[8] We provide a framework that can efficiently estimate individual effects for arbitrarily sized teams, and utilize it in the context of Congressional House personal staffs, where teams can contain up to 18 staffers. We demonstrate the advantages of our method in estimating a nested CES model, which is useful for understanding the optimal structure of teams. Our framework is suitable for any team environment where individual output is unobserved, such as in healthcare, innovation, and the workforce at large.

Third, this study contributes to the broader question of how political appointees and non-elected policymakers influence policy. A central tension in democracies is that while politicians are elected, appointees and bureaucrats are not, but they may nonetheless shape policy (Aberbach et al., 1981). Political appointees are influenced by partisan cycles (Spenkuch et al., 2023), and efforts to reduce the partisan influence on bureaucrats have increased their productivity (Aneja and Xu, Forthcoming). A broader discussion has centered on whether polarization is driven by the masses or elites (Enders, 2021; Boxell et al., 2024). We contribute to this literature by documenting how staffers may bridge ideological divides and especially aid in the creation of less-partisan legislation, potentially because they are unelected and relatively insulated from mass polarization. A rich literature has studied the structure and incentives of bureaucracies theoretically (Niskanen, 1968; Besley and Persson, 2009; Gailmard and Patty, 2012) and empirically (see Vogler (2022); Dahlström and Lapuente (2022) for two reviews). We estimate a production function of legislative output involving both politicians and a team of appointees.

Section 2 provides additional background on personal staffers in the US House and the legislative process. Section 3 lays out the datasets used. Section 4 contains a framework

[7] See Finkelstein et al. (2021); Abowd et al. (2008); Card et al. (2013); Bender et al. (2018); Song et al. (2019); Cantoni and Pons (2022); Chetty and Hendren (2018), among many others.

[8] Also related are Freund (2024), Iranzo et al. (2008), and Weidmann et al. (2024), who directly estimate individual effects and take them into team environments.

for estimation in team settings where individual contributions are unobserved, and is self-contained for readers interested in applying this framework. Section 5 estimates the staffer share of differences using a linear decomposition, a dynamic specification, and Bayesian estimation. Section 6 extends the Bayesian analysis to estimate sources of staffer productivity, examining the role that job titles, diversity, and within-team interactions play in shaping staffers' legislative impact. Section 7 estimates the ideology of staffers and discusses its implications for polarization and staffer productivity. Section 8 concludes.

2 Background

2.1 Personal staffers in the House

The modern US Congress employs over 17,000 staffers each year, spread across the many functions of Congress.[9] In this paper, we focus on the personal staff of Representatives of the House. These personal staff perform a variety of functions for their office: they respond to constituent concerns by phone or (e-)mail, they write speeches and other public-facing communication, they meet with lobbyists and organize their Representative's schedule, and they write bills that may someday become law.

Staffer turnover is relatively high. The hours tend to be long, and pay relatively low. In Table 1, we show that the average staffer is compensated at just under $45,000 per year. Given the difficulties of the job, staffers often exit Congress after a few years of work. The average staffer in our sample stays in Congress for just over 2 years, and often leaves the Hill for lobbyist jobs. Given the high turnover in staffers, Representatives are constantly on the search for new staffers to hire and fill old roles. Priority in hiring is typically given to anybody with experience on the Hill, with private mailing lists distributing job opportunities first before the general public is notified of job openings.[10]

Offices in Congress tend to be highly hierarchical. We show a typical org chart for a Representative's office in Figure A.1. Immediately reporting to the Representative themselves is the Chief of Staff, who in turn leads Directors in multiple domains (such as Legislation, Communication, or Outreach). The Legislative Director (LD) typically has Legislative Assistants or Aides (LAs) working for them, while Legislative Correspon-

[9] These include personal staff to members of Congress in both the House and Senate, committee staff in both the House and Senate, agency staff for the Congressional Budget Office, Government Accountability Office, and Congressional Research Service, as well as capitol police, legislative clerks, leadership staff, and others. See Cross and Gluck (2020) for a review of bureaucrats in Congress.

[10] See, for instance, `https://rollcall.com/2014/10/07/secrets-from-capitol-hills-back-rooms-how-to-get-hired-on-the-hill-2/`. Even entry into a full-time role in Congress is often predicated on securing an internship at Congress first.

dents (LCs) perform the lowest level of work. Oftentimes, LCs respond to constituent mail, though they may also conduct background research or assist with drafting legislation. Promotion up the job ladder often takes place within a year or two of beginning the job; in our data, the average staffer holds a job title for 1.3 years.

Each Representative is allowed to hire a maximum of 18 full-time staffers on a fixed budget (known as a Member Representational Allowance). Combined with the strict hierarchy of job titles, this can make promotion within an office difficult (if not impossible) if the staffer occupying the job above does not leave; thus, staffers often seek opportunities in other offices for promotion. Although ideology may play some role in these moving decisions, conditional on a Representative's political party, many staffers note that opportunities for upward mobility represent the most important reason for switching offices.[11]

The Legislative Process We provide a very brief summary of the legislative process in the US Congress, focusing on bills that originate in the House.[12] Before a bill is introduced to Congress, it is typically written by legislative staffers from one or more offices.[13] Upon introducing the bill to Congress, the Representatives of these offices become known as the "original co-sponsors" of the bill. The bill is then assigned to a committee, where it may be scheduled for markup (amendments), debates, and hearings, before eventually receiving a committee vote. Most bills die in committee. However, if a bill makes it through committee, it moves to the House floor where it is scheduled for debate and a roll call vote. A bill that passes the House roll call vote is sent to the Senate to go through a similar procedure.[14] If a bill passes both chambers of Congress, it is sent to the President, who may either sign the bill into law or veto it.

2.2 Some hypotheses

The power of Congressional staffers has been the subject of an extensive debate. Nonetheless, it has been long established that staffers are typically the primary authors of bills, conducting policy research, drafting legislative text, and coordinating with stakeholders to build coalitions (Fox and Hammond, 1977). A staffer's impact at this initial stage of

[11]This has been described to us in multiple interviews with former staffers.

[12]For a more detailed and digestible version of the process, see `https://www.congress.gov/legislative-process`. This process is explained in depth in Smith et al. (2013).

[13]On occasion, bills are also given to staffers by lobbyists. These may be modified before introduction to Congress.

[14]If the bill was first introduced in the Senate and passed the vote there, then it will be introduced to the House with sponsors (but not original co-sponsors), and follow the procedure outlined above. Bills may also be simultaneously introduced to both chambers of Congress, in which case a bill will still have an original co-sponsor in the House.

bill writing may thus depend on several factors: a staffer's job title, as legislative staff are directly involved in bill drafting while those in constituent services or office management may be more distant to the production of bills; the composition of their team, as diverse perspectives may improve legislation quality (Ritchie and You, 2021); and the effectiveness of their Representative, which may shape their own productivity (Hagedorn, 2015). However, during committee consideration, personal staff may play a reduced role as committee staff take the lead in managing hearings and markup sessions (Ommundsen, 2023). For floor votes, party leadership offices play a central coordinating role, with Representatives often relying on leadership for guidance rather than their personal staff. This institutional structure suggests that personal staffers may have their largest impact during the initial drafting of legislation.

While Representatives face strong electoral pressures that may constrain the ideological content of the bills they write (Mayhew, 2004), Congressional staffers are relatively insulated from the demands of voters. Some Representatives have explicitly stated that they do not consider the ideology of the staffers that they hire (Fox and Hammond, 1977). This puts staffers in a unique position, where they can influence the ideological content of legislation without being directly beholden to constituent preferences. It is ex-ante ambiguous whether this may lead them to polarize more than Representatives, or be a force towards moderation. Scholars have also found that the degree of legislative gridlock depends on the specific issue considered, with partisan conflicts contributing to gridlock (Flynn and Harbridge, 2016). It is thus also possible that the contribution of staffers to the legislative process may depend on the degree to which they are partisan.

3 Data

3.1 Statements of Disbursement

The primary dataset that we rely on is the House's Statements of Disbursement (SoD), which record all spending conducted by the House. Critically, these include payroll records that we use to identify the staffers who work for each office at a given point in time. The SoD has been published electronically on a quarterly basis since 2016. We utilize complete data on staffers from 2016 to 2022.

From 2.9 million payments made by Congress, we extract a total of 732 representatives and roughly 15,000 staffers.[15] We identify roughly 2,800 staffers that switch offices

[15]We disambiguate names using LinkTransformer (Arora and Dell, 2023). This is manually validated in the case of all Representatives; accuracy is generally high (over 99%).

("movers") in our main sample and construct panel data, associating a staffer with a single office in each quarter. A detailed description of the data construction process is outlined in Appendix A.

We code the job title of each staffer in each quarter following the keyword coding procedure in Crosson et al. (2021). We classify jobs into five categories: legislative staff, political management, communications, office management, and constituency service. We also code several job titles relevant to the legislative process: Chief of Staff, Legislative Director, Legislative Assistant/aide, and Legislative Correspondent.

We report summary statistics in Table 1. We find that roughly one-quarter of staffers work in a legislative role at any given point in time. In general, movers are not a representative sample of all staffers. They are disproportionately more likely to have longer careers, higher wages, and hold legislative or management positions. Given that we are primarily interested in understanding the importance of staffers on the whole relative to members of Congress, rather than in the effectiveness of the average staffer, we view this as a feature rather than a bug. The fact that movers are likelier to hold important positions in an office increases our ability to recover the true impact of staffers. However, the fact that we exclude a substantial fraction of staffers from the main analysis means that we are likely to underestimate the true effect of staffers.[16]

3.2 Other data sources

We supplement the House SoD with other data sources on Congress. First, we collect bill activity from the Congress.gov API. We use this data to construct our primary measures of legislative productivity: the number of original co-sponsored bills by each Representative in each quarter, the number of these co-sponsored bills that exit committee, the number of these co-sponsored bills that pass a House floor vote, and the number of these co-sponsored bills that are signed into law.

Second, we collect data on "important" bills from the Center on Effective Lawmaking, who define important bills as bills that receive an end-of-the-year write-up in the Congressional Quarterly Almanac.[17]

Finally, we collect data on roll call votes and DW-Nominate scores from `voteview.com` and election data from the Stanford-MIT Elections Performance Central project. We also collect staffer-level covariates from LegiStorm, including gender, education, and work experience, and impute ethnicity from names in a procedure described in Appendix A.

[16] In our Bayesian analysis, we are able to bring the non-movers back into the analysis and confirm this.

[17] As robustness, we also make use of the Legislative Effectiveness Scores that they produce. More details on this measure are provided in Appendix A.

Table 1: Summary statistics

	Overall			Movers	Non-movers	t-test
	Mean (SD)	Min	Max	Mean	Mean	p: 4 = 5
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Staffer level						
Career length (quarters)	8.84	0.00	38.00	14.93	7.84	0.000
	(7.76)			(7.89)	(7.27)	
Wage (dollars per quarter)	11193.11	0.00	52911.11	11823.19	11089.79	0.000
	(6424.67)			(5397.68)	(6572.17)	
Share career as legislative staff	0.23	0.00	1.00	0.33	0.21	0.000
	(0.39)			(0.41)	(0.39)	
Share career as managment	0.19	0.00	1.00	0.22	0.19	0.000
	(0.35)			(0.34)	(0.35)	
Share career as legislative director	0.03	0.00	1.00	0.07	0.03	0.000
	(0.15)			(0.18)	(0.15)	
Share career as chief of staff	0.04	0.00	1.00	0.05	0.04	0.005
	(0.19)			(0.20)	(0.19)	
Panel B: Staffer level (at least one year)						
Job length (quarters)	5.14	0.11	41.43	4.28	5.47	0.000
	(5.25)			(4.19)	(5.57)	
Office length (quarters)	6.89	0.11	41.43	5.63	7.40	0.000
	(6.68)			(5.35)	(7.09)	
# jobs (per office)	1.34	1.00	7.00	1.32	1.35	0.000
	(0.68)			(0.65)	(0.69)	
Ever switch job category (in office)	0.88	0.00	1.00	0.88	0.88	0.844
	(0.33)			(0.33)	(0.33)	
# offices over career	1.36	1.00	132.00	2.74	1.13	0.000
	(2.67)			(4.76)	(2.04)	

Notes: This table presents summary statistics for personal staffers in the House of represents. In Panel A, we present summary statistics for all staffers. Wages are in US dollars per quarter. In Panel B, we present summary statistics for all staffers who work in the House for at least one year. Job length is the amount of time (in quarters) that a staffer holds the same job title. Office length is the amount of time (in quarters) that a staffer works for the same Representative. Standard deviations presented below means in parenthesis.

4 A Mover Design for Team Environments

In this section, we present a framework for estimating the staffer share of differences between the offices of Representatives in Congress.

4.1 Linear decomposition

In the canonical mover design, an individual level outcome such as wages (Abowd et al., 1999), mortality (Finkelstein et al., 2021), or voting behavior (Cantoni and Pons, 2022) is estimated using a linear combination of individual fixed effects and group (often location)

fixed effects:

$$y_{ijt} = \alpha_i + \gamma_j + \tau_t + \varepsilon_{ijt}$$

where i is an individual, j is a group, and t is a time period. In our setting, i is a staffer, j is a Representative's office, and t is a quarter.

Unfortunately, there are many circumstances in which group-level outcomes are influenced by individuals, but the specific contribution of each individual to the group is unknown, making this equation impossible to estimate. For instance, in our setting, we do not observe the specific legislation that each staffer produces. However, given that we observe the movements of staffers across different offices, as well as the legislative output of offices during these moves, we can exploit the information contained in these moves to inform our estimates of each staffer's productivity. We extend the standard mover design to team environments, continuing to assume a linear value-added model, or production function, of legislation. We relax this linearity assumption in Section 6 and estimate a nested CES model, among others.

Let $m_{i,j,t} = \mathbb{I}[\text{staffer } i \text{ in office } j \text{ at time } t]$. Then our model is:

$$y_{jt} = [\sum_i \alpha_i m_{i,j,t}] + \gamma_j + \tau_t + \varepsilon_{jt} \tag{1}$$

where m and y_{jt} are data, and the remaining terms are parameters to be estimated. We assume that $E[\varepsilon_{jt}|\alpha_i, \gamma_j, \tau_t, m_{ijt}] = 0$.

Our estimand of interest is the share of differences in legislative output between high and low productivity offices that is attributable to staffers. Under the standard AKM mover design, this can be identified via a linear decomposition. A similar argument holds when we estimate the team mover design in Equation 1.

Let $\bar{y}_j = \gamma_j + \bar{y}_j^{\text{staff}}$ be the expectation of y_{jt} over all time periods, where

$$\bar{y}_j^{\text{staff}} \equiv \frac{1}{T} \sum_t \sum_i \alpha_i m_{i,j,t}$$

and we impose $\frac{1}{T}\sum_t \tau_t = 0$. Hence, for any two offices d, o:

$$\bar{y}_d - \bar{y}_o = (\gamma_d - \gamma_o) + (\bar{y}_d^{\text{staff}} - \bar{y}_o^{\text{staff}}) \tag{2}$$

Equation 2 shows that the difference in average office output between any two offices can be decomposed into two parts: the first part is the difference between the office fixed effects $\gamma_d - \gamma_o$, and the second part is the difference between staffer-specific components $\bar{y}_d^{\text{staff}} - \bar{y}_o^{\text{staff}}$. The shares of differences in office output between origin and destination offices that are driven by staffers and offices (including Representatives) are:

$$S^{\text{office}}(d, o) = \frac{\gamma_d - \gamma_o}{\bar{y}_d - \bar{y}_o},$$

$$S^{\text{staff}}(d,o) = \frac{\bar{y}_d^{\text{staff}} - \bar{y}_o^{\text{staff}}}{\bar{y}_d - \bar{y}_o} = 1 - S^{\text{office}}(d,o)$$

This allows us to estimate the staffer share of differences in output between any pair of offices. Throughout the paper, we estimate S^{office} and subtract it from 1 for a consistent estimate of our parameter of interest: the staffer share S^{staff}.

We focus on estimating the difference between two sets of offices O and D—most frequently, the difference between above and below median offices in bills produced. For an office level quantity x_d, define $x_D = \frac{1}{|D|}\sum_{d\in D}\bar{x}_d$. Then the office and staffer share of differences is simply:

$$S^{\text{office}}(D,O) = \frac{\gamma_D - \gamma_O}{\bar{y}_D - \bar{y}_O} = 1 - S^{\text{staff}}(D,O) \tag{3}$$

While we term S^{staff} and S^{office} the staffer and office share respectively, in reality, they capture all forces that are always associated with the staffer or the Representative's office. Thus, if there are non-mover staffers who are always attached to the same Representative throughout our sample, we will attribute this staffer's impact to the office share, and hence likely underestimate the true staffer share. A Representative's office may also appear effective not because of the Representative's inherent quality, but because of external factors (for instance, they happen to be a committee chair for the entirety of our sample, or they represent an important constituency). Similarly, lobbyists and other actors may follow staffers rather than Representatives, providing staffers with policy ideas or literal bills over the course of their career; the impact of these people will be bundled into the estimated staffer fixed effects. As such, the staffer share that we compute should be interpreted as representing the staffers' contribution broadly construed, inclusive of all connections or other permanent factors associated with them.

4.2 Two dynamic estimating equations

Estimating staffer shares dynamically In order to track changes in office productivity around the timing of a staffer move, we estimate an event-study version of Equation 1. To do so, we extend the derivation in Cantoni and Pons (2022) to a team setting. We first rewrite Equation 1 as follows:

$$y_{jt} = \alpha_i m_{ijt} + [\sum_{i'\neq i} \alpha_{i'} m_{i'jt}] + \gamma_j + \tau_t + \epsilon_{jt} \tag{4}$$

Consider a staffer i who moves from an origin office $o(i)$ to a destination office $d(i)$. Let $r(i,t) = t - t_i^*$ be the time to a staffer i's first move and $\delta_i \equiv \bar{y}_{d(i)} - \bar{y}_{o(i)}$ be the difference in average office output between the destination and origin office for staffer i. By

substitution, Equation 4 can be rewritten as:

$$y_{jt} = \alpha_i m_{ijt} + \gamma_{o(i)} + \mathbb{I}_{r(i,t)\geq 0} \times S^{\text{office}}(d(i), o(i)) \times \delta_i + [\sum_{i' \neq i} \alpha_{i'} m_{i'jt}] + \tau_t + \epsilon_{ijt}$$

Let $\tilde{\alpha}_i = \alpha_i + \gamma_{o(i)}$, substitute $\mathbb{I}_{r(i,t)\geq 0}$ with indicators for time to move, and substitute δ_i with its sample analogue $\hat{\delta}_i = \hat{y}_{d(i)} - \hat{y}_{o(i)}$. The event-study specification becomes:

$$y_{jt} = \tilde{\alpha}_i m_{ijt} + \theta_{r(i,t)} \hat{\delta}_i + [\sum_{i' \neq i} \alpha_{i'} m_{i'jt}] + \tau_t + \epsilon_{jt} \tag{5}$$

where our estimand of interest is $\theta_{r(i,t)}$, the dynamic measure of the office share of differences at $r(i,t)$. This estimate is useful because it allows us to understand how quickly a staffer impacts an office once they move, whether their effect persists, and to evaluate pre-trends in office behavior. We note that because staffers may move from any origin to any destination office, the value of θ after a staffer moves may differ from the staffer share of differences between high and low productivity offices, but we expect these quantities to be similar if there is no systematic staffer sorting through moves.

Equation 5 is computationally expensive to estimate, so we recast it in a form that is estimable by standard fixed effect routines.[18] Let $\hat{\alpha}_i$ be an unbiased estimate of α_i. This motivates the use of the following outcome:

$$\hat{y}_{it} = \mathbb{I}_{r(i,t)\geq 0}[y_{d(i)t} - \sum_{i' \neq i} \hat{\alpha}_{i'} m_{i',d(i),t}] + \mathbb{I}_{r(i,t)<0}[y_{o(i)t} - \sum_{i' \neq i} \hat{\alpha}_{i'} m_{i',o(i),t}]$$

and event study specification:

$$\hat{y}_{it} = \tilde{\alpha}_i + \theta_{r(i,t)} \hat{\delta}_i + \tau_t + \epsilon_{it} \tag{6}$$

Notice that Equation 6 is equivalent to the event study specification in Equation 5, provided that $\alpha_i = \hat{\alpha}_i$.[19] It has the advantage of being estimable by standard fixed effect routines as we no longer have a variable number of fixed effects per observation to estimate. However, it requires $\hat{\alpha}_i$ to be an unbiased estimate of α_i. Reusing $\hat{\alpha}_i$ from the same data is problematic as it uses the data twice to validate the estimate. Instead, we split our sample into odd quarters and even quarters. We first estimate Equation 1 above on odd quarters only. Then, we turn to even quarters using the fixed effects estimates from odd quarters to estimate Equation 6.

[18] As discussed in Appendix B, we estimate Equation 1 using the full design matrix containing indicators for every staffer and $O(JT)$ rows of data, where J is the number of offices and T the number of time periods. However, once we move to an event study specification, we have $O(IT)$ rows of data, where I is the number of staffers, and $I >> J$.

[19] In the case where $\alpha_i = \hat{\alpha}_i$, the errors $\epsilon_{it} = \mathbb{I}_{r(i,t)\geq 0}\epsilon_{d(i),t} + \mathbb{I}_{r(i,t)<0}\epsilon_{o(i),t}$. This motivates clustering errors at the office level. Our results are also robust to the inclusion of time fixed effects relative to the move $\rho_{r(i,t)}$.

Predicting a staffer's movements We continue to analyze a staffer's first move between origin office o and destination office d. Our goal is to assess whether we can predict the timing of a staffer's move by examining the correlation between their out-of-sample fixed effect and the legislative output of their destination office. If successful, this approach validates the estimated staffer and office fixed effects

We develop a new estimation strategy to validate Equation 1 in a dynamic fashion. Specifically, we estimate:

$$y_{d(i),t} = \sum_{\tau=-4}^{4} \mathbb{I}[r(i,t-\tau)=0] \left(\theta_\tau \hat{\alpha}_i + \eta_\tau \left(\sum_{i' \neq i} \hat{\alpha}_{i'} m_{i',d(i),t} \right) + \zeta_\tau \hat{\gamma}_{d(i)} \right) + \hat{\tau}_t + \varepsilon_{d(i),t} \quad (7)$$

where $\theta_\tau \hat{\alpha}_i$ captures the specific contribution of a staffer i to the destination office's output at each time lag τ, $\eta_\tau \left(\sum_{i' \neq i} \hat{\alpha}_{i'} m_{i',d(i),t} \right)$ captures the aggregate effect of other staffers on the destination office's output at each time lag τ, and $\zeta_\tau \hat{\gamma}_{d(i)}$ captures the time-varying Representative-level fixed effect at each time lag τ.[20]

If the model in Equation 1 is true, then we should expect the following to hold: (1) θ_τ should be 0 prior to the move (because the staffer is not in the destination office), and positive after the move, (2) η_τ should be constant, because the team of staffers at the destination office is always present at the destination, and (3) ζ_τ should be constant, because the destination office Representative is always in the destination office. Note that if these conditions hold, then Equation 7 directly collapses to Equation 1, except stacked at the individual level.[21]

Estimating this equation serves as a validation of the model in three ways: (1) a $\theta_\tau > 0$ means that we are able to predict when a staffer has moved based purely on their fixed effect. Since $\hat{\alpha}_i$ is estimated out of sample, this signals that an individual fixed effect is estimated with relative precision, (2) an η_τ and ζ_τ that remains constant over time serves as both a placebo test and a validation that these fixed effects are well estimated, and (3) pre-trends can be tested across θ_τ, η_τ, and ζ_τ for $\tau < 0$. Thus, we can test for anticipatory behavior in the staffer, their team, and the Representative all at once. Finally, we note that by replacing the destination office with the origin office in Equation 7, one can examine whether staffers have a persistent effect on offices even after they leave.

[20] We face the same problem as the event study estimation if the same data is used to estimate fixed effects α_i, γ_j as well as the main event study. Thus, we estimate these fixed effects on odd quarters and estimate Equation 7 in even quarters.

[21] This once again motivates clustering standard errors at the office level.

4.3 Bayesian estimation

We next move to estimating the staffer share of differences using random effects. This is primarily so that we can estimate richer, nonlinear models of team production and staffer ideology. With fixed effects, estimating a large number of parameters leads to issues with consistency. Our data-generating process is such that the number of parameters (individual fixed effects) grows with the size of the dataset, but we do not gain more information about individuals as the dataset grows—collecting more years of data does not help when an individual has already left Congress. As the number of parameters increases, the errors in these estimates do not necessarily average out, leading to biased or inconsistent parameter estimates (Andrews et al., 2008). This is particularly problematic in nonlinear models, where this can lead to other parameters of interest being biased.

On the other hand, by modeling α_i and γ_j probabilistically, we overcome the incidental parameters problem discussed above.[22] Our model is Equation 1 such that:

$$y_{jt}|\alpha_i, \gamma_j, \tau_t, m_{ijt} \sim \mathcal{N}([\sum_i \alpha_i m_{ijt}] + \gamma_j + \tau_t, \sigma^2) \tag{8}$$

$$\alpha_i \sim \log\mathcal{N}(\mu_\alpha, \sigma_\alpha^2), \quad \gamma_j \sim \log\mathcal{N}(\mu_\gamma, \sigma_\gamma^2), \quad \tau_t \sim \log\mathcal{N}(\mu_\tau, \sigma_\tau^2)$$

with priors for $k \in \{\alpha, \gamma, \tau\} : \mu_k \sim \mathcal{N}(0, \sigma_\mu), \sigma_k \sim \mathcal{N}(0, \sigma_\sigma)$. We estimate this model using Hamiltonian Monte Carlo (HMC), a type of Markov Chain Monte Carlo technique. This method makes the estimation of problems with difficult geometries and large numbers of parameters tractable.[23] With estimates of these random effects in hand, we can follow the procedure outlined in the sections above to derive the staffer share of differences.[24]

Although we depart from the empirical Bayes approach used by many economists in favor of a pure Bayesian approach,[25] we note that (under regularity conditions) the Bernstein-von Mises Theorem means that asymptotically, Bayesian credible intervals are equivalent to frequentist confidence intervals and is efficient. The primary advantage of our approach is that the Bayesian theory of inference holds in finite samples. Using this approach, we can evaluate how staffers' job titles influence their legislative productivity, incorporate a nested CES structure to analyze how staffers and Representatives might substitute or complement one another, and study how staffers and Representatives affect

[22] This also helps deal with issues like overfitting and multicollinearity that can arise when the number of parameters to be estimated is large relative to the sample size. For more work on correlated random effects in the mover design, see for instance Bonhomme et al. (2019).

[23] We describe HMC in more detail in Appendix B. To the best of our knowledge, the only other paper in economics that makes use of HMC is Childers et al. (2022).

[24] We leave a discussion of choice of priors to Section 5.

[25] For some examples of empirical Bayes used in value-added models, see examples such as Chetty et al. (2014); Angrist et al. (2017); Chetty and Hendren (2018). Purely Bayesian approaches have also been used before, as in Abowd et al. (2015).

the ideological content of bills.[26]

4.4 Discussion

In a recent review of personnel economics, Hoffman and Stanton (2024) note that while "understanding teams and team production is of huge practical importance," the actual literature on teams is "relatively sparse." In Appendix B, we compare our approach to other leading proposals in the literature, focusing on performance in our context, where teams may contain up to 18 staffers.[27] Most procedures turn out to be infeasible to estimate, because the parameter space grows sufficiently quickly in team size that computation is not possible or parameters are no longer identified. In Section 6.3, we compare our method against Ahmadpoor and Jones (2019), who provide a feasible method in our context, and show that our method has advantages. Researchers have also adopted misspecified models that fail to correctly account for the impact of teams. This typically involves conducting analysis at the individual level or only focusing on one key individual per team.[28] In Section 5, we discuss how such misspecification would change our findings, leading to estimates half as large as the correctly specified model.

Thus, although we focus on estimating the effect of Congressional staffers in this paper, we believe that our approach to estimating team effects holds some advantages over prior work. In domains ranging from healthcare to innovation to finance to the workforce at large, there are many team environments where individual contributions are unobserved and this framework may be applicable.

5 The Staffer Share of Legislative Productivity

How important are staffers? We estimate the share of differences in legislative productivity between offices that can be attributed to staffers. To estimate this quantity, we employ three distinct approaches: static linear decompositions, dynamic estimates of the staffer share, and a Bayesian approach. Each method provides unique insight into the role staffers play in producing legislation, and the similarity in findings across approaches may build confidence in the overall result.

[26] We discuss specific model extensions in more detail in Sections 6-7.

[27] In our entire sample, the average team contains 12.94 staffers. In the mover sample, the average team contains 4.91 staffers. With the Bayesian approach, we are able to make use of the full team.

[28] Constantine and Correia (2021) list examples of papers that are misspecified in this way. They also provide an implementation of the linear model in Equation 1, but because their procedure treats the FEs as nuisance parameters that are not directly estimated, it cannot be used for decompositions.

5.1 Linear decompositions

We begin by estimating a series of linear decompositions following Equation 1 of our mover sample. This allows us to estimate staffer effects (α_i) and office effects (γ_j) for each staffer and office in the sample. In each quarter, we identify the above and below median offices for five legislative outcomes of interest: the number of original co-sponsored bills introduced by the office, the number of these bills that exited committee, the number of these bills that passed a House floor vote, the number of these bills that were signed into law, and the number of "important" bills introduced.

We use these estimated effects and Equation 3 to obtain a consistent estimate of the staffer share of differences between high and low output offices, S^{staff}. These staffer shares are plotted in the blue bars of Figure 1.[29]

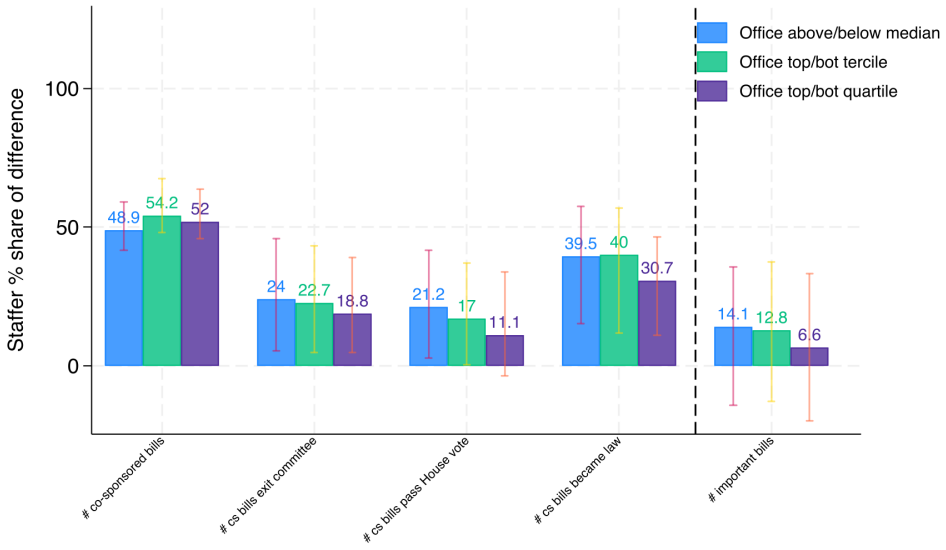
We find that staffers are a critical input to the most basic output of Congress: bills and laws. 49% of the difference in bills produced between above and below median offices can be attributed to staffers. As we move across the life-cycle of a bill, we find a U-shaped pattern for the importance of staffers. Staffers account for only 20 to 25% of the difference in the number of co-sponsored bills that exit committee or pass a House floor vote. Given that each committee has its own staff who handle the text of bills, and the importance of political considerations in ushering a bill out of committee and through a House vote, we find the smaller role of personal staffers in these intermediate stages of a bill's life-cycle to be plausible. However, the importance of staffers appears to tick up again, to roughly 40%, once we examine the share of differences in co-sponsored bills that eventually become law. After the passage of a bill through the House, it must also pass the Senate and be signed by the President. Inter-branch negotiation and coordination may require more manpower and connections, both of which personal staff may provide.

We also examine the staffer share of differences in the number of important bills co-sponsored by the office. We find that staffers play a much smaller role in the introduction of these important bills, with a point estimate of 14.1% and a staffer share of 0 falling within the 95% confidence interval. We attribute this to several potential factors: first, this is a low frequency outcome. While the average office co-sponsors 16.3 bills in a quarter, only 0.2% of these bills are classified as important (for reference, 3.4% of bills in our sample are signed into law). Second, bills are classified as "important" if they receive a year-end write-up in the Congressional Quarterly Almanac; a bill may be likelier to receive such a write-up if the Representative is spearheading the bill and promoting or

[29] We form confidence intervals using the Bayesian bootstrap. More details on the exact estimation procedure are contained in Appendix B.1.

advertising it, compared to bills originated from the staff. Finally, we note that bills classified as "important" are likely to be the subject of intense political dispute. An effective bureaucracy may focus on issues of policy rather than politics, which may make such bills less contentious and worthy of a write-up.

Figure 1: Staffer share of differences in legislative productivity, linear decomposition



Note: This figure plots the staffer share of differences between offices that are above and below median in legislative output in blue, as well as 95% confidence intervals. The staffer share is computed from fixed effects estimated in a team-based mover design framework (Equation 1). Staffer shares for the difference in top and bottom tercile (quartile) offices are plotted in green (purple). Bars are arranged from left to right in order of closeness to passage into law, except for 'important bills', which refers to the number of bills referenced in the Congressional Quarterly Almanac year-end write-up. 'cs bills' refers to the number of bills originally co-sponsored by the office.

Robustness and IV For robustness, we also compute the staffer share of the difference between top tercile and bottom tercile offices (plotted in green) and between top quartile and bottom quartile (plotted in purple). We find that our estimates are very similar, even when looking at offices with more more extreme differences in output.

We perform two validations of the staffer fixed effect estimates. First, we regress the absolute value of each staffer's fixed effect $|\alpha_i|$, to capture the impact of a staffer (good or bad), against their average wage w_i in Table 2, Panel A. We find that a standard deviation increase in a staffer's impact is associated with a 5 to 10% standard deviation increase in wages, consistent across outcomes. An attentive reader may notice that this regression coefficient is not identified: the fixed effects of two staffers with identical moving histories

cannot be separately identified.[30] To address this, in Panel B, we collapse the data to the work history level (taking an average across wages) and find similar results.[31] To demonstrate that this correlation is not driven by the difference between major job titles, in Panels C and D, we focus on the subset of staffers that are managers, or managers and legislative staff. We find that, if anything, the relationship between a staffer's impact and their wage is stronger within these samples.[32]

Table 2: Correlating staffer FEs with wages

FE source:	Bills cospons.	Laws passed	Exit comm.	Pass House	Important bills	Index
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: staffer level						
abs(Staffer FE)	0.080	0.056	0.114	0.119	0.042	0.084
	(0.023)	(0.022)	(0.025)	(0.025)	(0.018)	(0.018)
Panel B: history level						
abs(Staffer FE)	0.084	0.062	0.121	0.125	0.042	0.089
	(0.023)	(0.022)	(0.025)	(0.025)	(0.018)	(0.018)
Panel C: history level, managers only						
abs(Staffer FE)	0.124	0.203	0.327	0.318	0.097	0.221
	(0.056)	(0.065)	(0.055)	(0.056)	(0.054)	(0.047)
Panel D: history level, managers and legislative staff only						
abs(Staffer FE)	0.099	0.101	0.169	0.170	0.048	0.123
	(0.033)	(0.033)	(0.036)	(0.036)	(0.031)	(0.027)

Notes: Regressions are at the staffer level in Panel A, and at the work history level in Panels B-D (independent and dependent variables for all staffers that are always in the same office at the same time are averaged to a single observation). In Panel C, we restrict the sample to staffers who work in a management position. In Panel D, we restrict the sample to staffers who work in a management or legislative position. Independent and dependent variables of interest have been standardized to mean 0, and standard deviation 1. Staffer fixed effects α_i are estimated from the linear team model in Equation 1. Each column presents an estimate using a different outcome variable, and column 6 is an index consisting of the sum of the prior five columns, standardized. Wages are a staffer's average wage over their career. Standard errors are robust.

Second, we correlate staffer fixed effects across specifications and present results in Table A.1. Despite the variation in staffer shares across outcomes, it appears that productive staffers are helpful across all stages of the bill life-cycle: the within-staffer correlations are consistently larger than 0.50 and statistically significant.

[30] While this regression coefficient is not identified, we emphasize that the staffer share of differences *is*, due to the linearity of the underlying decomposition. Once we move to nonlinear models, we use Bayesian estimation, where the distribution on α_i disciplines the staffer random effect estimates.

[31] We can also restrict attention to only staffers with unique work histories (over 95% of the sample) or weight histories by the number of staffers, with extremely similar results (available on request).

[32] Strictly speaking, this validation exercise jointly tests how well for our staffer effects α_i are estimated, as well as whether staffer wages are rigid or responsive to staffer productivity. The positive coefficient suggests that wages are connected to staffer productivity, an idea we return to in Section 6.

We also test for the staffer share of legislative productivity using the Center for Effective Lawmaking's Legislative Effectiveness Score (LES) and the LES minus a benchmark of Representative characteristics (including seniority, chair status in committee, and party in the majority). We present results in Figure A.2. We estimate that the staffer share of differences in legislative effectiveness is roughly 40% in both outcomes, consistent with our findings of the staffer share in bill production and passing laws.

We stress the importance of using the team-based design in Equation 1 by comparing it to misspecified approaches that fail to properly account for the presence of teams. In Table A.2, Panel A, we present our baseline estimates from the team-based mover design. Panel B presents estimates from the standard AKM design at the individual level. Across outcomes, the estimated staffer share is at most half as large as the baseline. Focusing on key individuals within the team does not help matters: in Panels C and D, we restrict attention to Chiefs of Staff and Legislative Directors respectively, and find that the staffer share of differences is similarly underestimated.

One may also be concerned about the exogeneity of staffer moves. Although qualitative work suggests that opportunities up the job ladder represent the dominant reason for moving between offices, there are many potential reasons as to why staffer productivity may be correlated with their movement.[33] We explicitly test for persistent or slow-moving sources of endogeneity in Section 5.2. This leaves the possibility of fast-moving sources of endogeneity. We address these using an instrumental variables strategy.

We deploy two sets of instruments for staffer movements between offices: first, we exploit the fact that following turnovers in Congress, staffers often move to other offices within the same state (a home-bias effect), and use an indicator for a staffer's origin office losing an election and the potential destination office being in the same state as an instrument. Second, we draw on the migration literature and interact a "push" factor between two offices (the leave-out total flow of staffers between the origin and destination) with a "pull" factor into the destination (the number of staffers who leave the destination in that period). Given the high number of moves that we are instrumenting for, a very strong first stage is required. To achieve this, we deploy machine learning methods to estimate the first stage.[34] We present results from this estimation procedure in Figure A.2. The IV estimate of the staffer share in co-sponsored bills is 85.5%, a substantially larger share

[33]For example, more productive staffers may be poached by more productive offices, while less productive staffers may be disproportionately fired by some offices; staffers may have ideological preferences in choosing offices to work for, which may be correlated with productivity; staffers may go to offices where they think they will be more impactful.

[34]The complete estimation procedure, as well as more details about the instrument, is contained in Appendix B.1.

than the OLS estimate.[35]

Finally, we raise one point of interpretation. Staffers and offices may differ in their legislative productivity due to differences in their objective functions (and not in their capabilities): this is most clearly seen in the case of non-legislative staffers, for whom we should not estimate high legislative productivities, even if they might have been very productive in legislative positions. Similarly, some Representatives may prioritize non-legislative functions of their job, such as constituent service, government oversight, or media relations, which would translate into lower estimated productivities. To the extent that this is true, the differences in office productivities that we attribute to staffers may reflect a combination of differences in capabilities and in preferences.

5.2 Dynamic estimation

We proceed by estimating two sets of dynamic specifications. This allows us to quantify how quickly a staffer begins to influence their office, as well as test for potential pre-trends and validate our linear decomposition estimates.

Dynamically estimating the office share of differences in legislative productivity We estimate event studies of the office share of differences in legislative productivity between a staffer's origin and destination office. Throughout our dynamic analysis, we focus on the first move that a staffer makes to simplify the interpretation of the estimated quantities. As noted in Section 4, we estimate all fixed effects on odd quarters and estimate the event study Equation 6 on even quarters. In Figure 2, we plot the estimated office share ($\hat{\theta}_{r(i,t)}$) in red and staffer share ($1 - \hat{\theta}_{r(i,t)}$) in blue, along with 95% confidence intervals with underlying standard errors clustered at the office level.

We find no evidence to support pre-trends: across all specifications, the estimated office share is never different from 0 with statistical significance at any time period. The difference in productivity between a staffer's origin and destination office are not correlated with the origin office's productivity $\hat{y}_{ijt}$ (which includes the staffer's own productivity α_i). This supports the causal interpretation that we advance, as it rules out slow-moving sources of endogeneity (such as a staffer becoming more productive over time, being noticed by a high-productivity office, and hired for that reason). Because the office share is close to 0, the staffer share is mechanically close to 1.

We note that there is a sharp increase (decrease) in the office (staffer) share in the exact quarter that the staffer moves in all specifications. This can be interpreted as the share

[35] Given the large standard errors in this estimate, we prefer to interpret the IV estimates as a sanity check that may help us to sign bias, rather than relying on the specific quantitative estimate.

of differences between the destination and origin offices that is attributable to the office (staffer). This share remains stable for the remaining quarters. In Figure 2, Panel A, we find that the estimated staffer share of differences in bills produced is typically around 40%, although the confidence intervals frequently include the 50% quantity estimated in the linear decomposition.[36] In Panels B-D, we find that the staffer share of differences through the bill life-cycle also hovers around the 40-50% range. In Panel E, examining important bills, we once again find that the staffer share is close to 0%.

The instant change at $t = 0$ and the relative stability of coefficients across quarters suggests that staffers immediately have an impact upon moving offices—it does not take multiple quarters for their presence to be felt. We additionally validate this finding with an in-sample estimation on the odd quarters. We present results in Figure A.3, noting the absence of pre-trends, the immediate jump in office shares in the quarter the staffer moves, and the stability of the coefficient in the periods afterward.

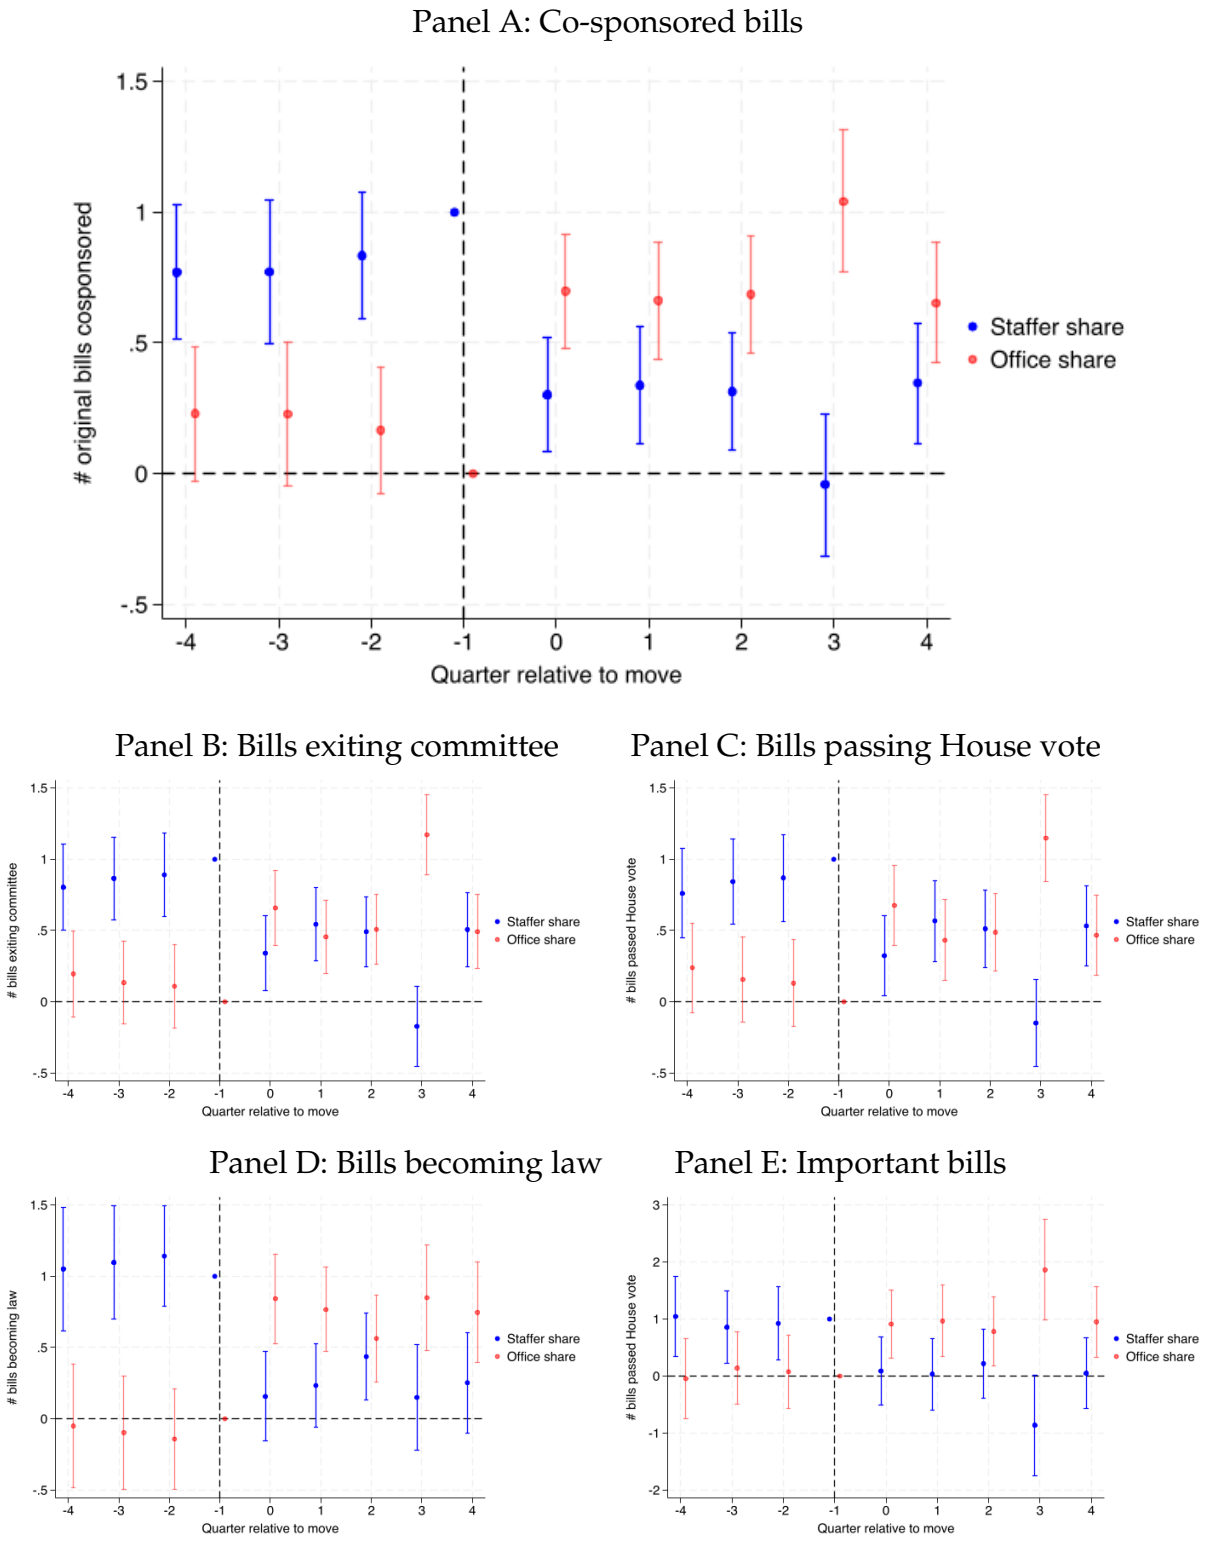
Taking these results on the whole, we once again find that staffers play a substantial role in the legislative process.

Dynamically predicting staffer movement Instead of estimating the office share of differences in legislative productivity, we can also predict the quarter in which staffers move offices. We take the estimated office and staffer fixed effects from odd quarters and estimate Equation 7 in even quarters, using bills produced as the outcome. In Figure 3, we plot the estimated coefficients on a staffer's own FE (θ_τ) in red, along with 95% confidence intervals with underlying standard errors clustered at the office level. Intuitively, we regress out-of-sample estimates of a staffer's productivity on the destination office's legislative output at varying time lags. Any positive coefficient means that we have detected the presence of a staffer at an office without using information on when they moved. We expect the coefficient to be 0 prior to the move (as they are not in the office), and to be positive after the move.

We first note the absence of pre-trends in a staffer's own fixed effect, once again supporting a causal interpretation. Prior to a staffer moving, there is no correlation between their own productivity $\hat{\alpha}_i$ and the destination office's output $y_{d(i),t}$. Strikingly, we are able to predict the presence of a staffer moving into an office using only their (out-of-sample) fixed effect, indicating that these are estimated with some precision. However, they are not measured perfectly. We plot the model implied coefficient when out-of-sample FEs

[36] We note that the underlying shares estimated are different quantities: in the linear decomposition, this is the share of differences between high and low-productivity offices, while in the dynamic specification, this is the share of differences between destination and origin offices. Staffers do not systematically move to higher productivity offices.

Figure 2: Dynamic estimation of the staffer share of differences in legislative productivity



Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 6) at the staffer-quarter level (for even quarters only). The x-axis shows quarters to a staffer's first move. Displayed coefficients in blue (red) are the staffer (office) share of differences between the destination and origin office's legislative output. The specific type of legislative output varies by panel. Fixed effects for other staffers on the team (estimated out of sample on odd quarters) are subtracted from the outcome. Standard errors are clustered at the office level.

are estimated without error in blue, which predicts that θ_τ should be exactly 1 in the post-period. Estimation error in our staffer fixed effects could explain the gap between our estimates (in red) and the benchmark (in blue). We conduct an in-sample validation on odd quarters in Figure A.5, Panel A. In-sample, we find that the own staffer coefficient θ_τ is now much closer to 1 in the post-period, supporting this idea.[37] Finally, we present coefficients from a placebo test where we randomly permute a staffer's destination office 100 times; we find that our pre-period estimates are comparable to the permuted estimates, while our post-period estimates always lie above them.

Our estimating equation, Equation 7, simultaneously also regresses out-of-sample estimates of the other staffers on the team and office effects on the outcome, which we expect to be constant as they are always present at the destination office. In Figure A.4, Panel A, we present coefficients on the office (ζ_τ) in blue and the other team members (η_τ) in gray. Following the theoretical prediction of the model, we find that these coefficients remain relatively stable over time, exhibiting only minor deviations from the normalized coefficient at $\tau = -1$ in both the pre- and post-period. This also lends empirical support to the model that we estimate.

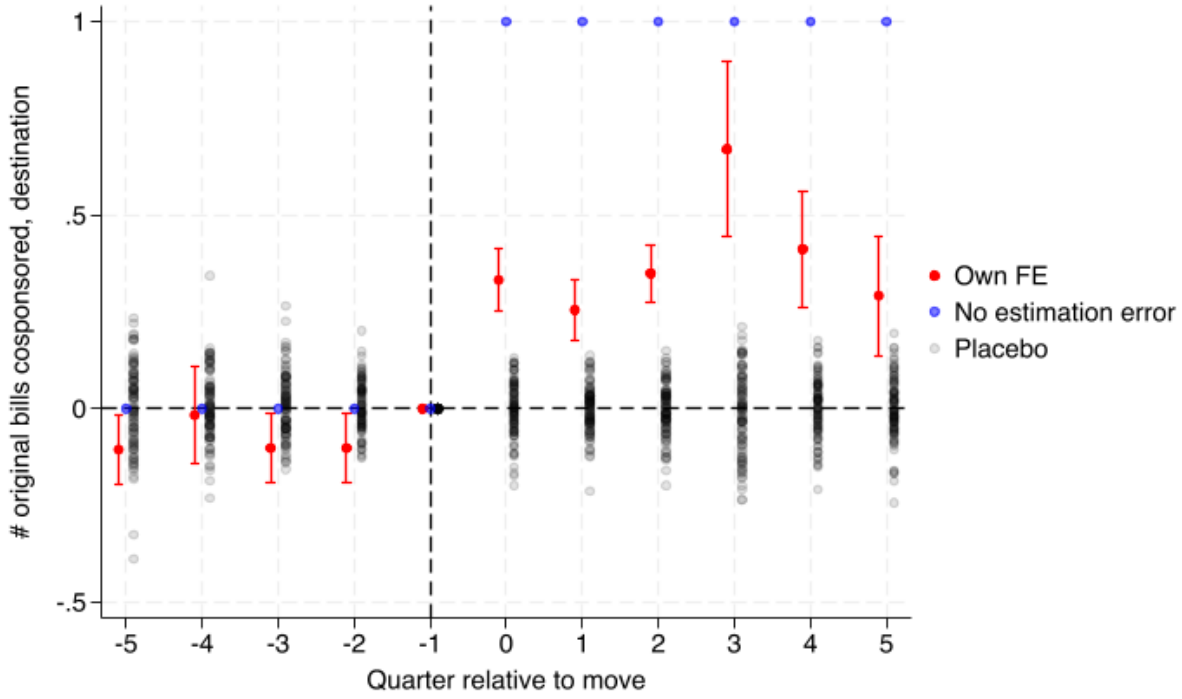
We present results using other measures of legislative productivity as the outcome in Figure A.4, Panels B-E. We are able to successfully predict the movement of staffers from their out-of-sample fixed effect when examining bills exiting committees and bills passing a House vote. Unfortunately, for the low-frequency events of bills becoming law and introducing important bills, we are not able to estimate staffer effects with enough precision to predict a staffer's move with statistical significance.[38]

Finally, we estimate whether staffers have a persistent effect on their offices after they leave. To do so, we re-estimate Equation 7, except we now use the output from the origin office $y_{o(i),t}$ as the outcome instead. We plot results out-of-sample and in-sample in Figures A.6-A.7. We find that staffers have no persistent effect on their offices: in the very quarter that they leave, we cannot distinguish their effect from 0, whereas the effect of the Representative and other staffers remains constant. Thus, the impact of staffers on legislation appears to be contingent on their physical presence, rather than any deeper structural or cultural changes they may impart on an office.

[37] Our post-period estimates may also be less than 1 due to model misspecification—assuming linearity across all staffers and the Representative is a strong assumption, and one that we explicitly test in Section 6.

[38] In Panel E, a joint test of the post-period own staffer fixed effects being different from 0 has a p-value of 0.103. Using the in-sample fixed effects in Figure A.5, Panels B-E, we are always able to predict a staffer's move, regardless of outcome.

Figure 3: Out of sample correlation between a moving staffer's fixed effect and destination bills co-sponsored



Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 7) at the staffer-quarter level (for even quarters only). The x-axis shows quarters to a staffer's first move. The outcome is the number of co-sponsored bills produced by the destination office. The full regression includes staffer fixed effects, office fixed effects, and fixed effects for the rest of the team. Coefficients on the moving staffer's fixed effect are plotted in red. Coefficients for an ideal benchmark under which out of sample staffer FEs are estimated without error are plotted in blue. Coefficients for a placebo exercise, where the destination office is randomly permuted 100 times, are plotted in gray. All fixed effects are estimated out of sample on odd quarters. Standard errors are clustered at the office level.

5.3 Bayesian estimation

We next estimate our original team-based mover design model, specified in Equation 1, using Bayesian methods. We employ fully Bayesian methods for two reasons: first, it provides a finite sample theory of inference for reasons discussed in Section 4. Second, to estimate richer, non-linear models of staffer productivity and ideology, we will need credible estimates of staffer effects. Unfortunately, we cannot guarantee consistency of fixed effects (Neyman and Scott, 1948)—we find evidence of estimation error in our fixed effects, which will not average out in nonlinear models. However, using a hierarchical model with random effects, we can consistently and efficiently estimate the distribution of these effects. Thus, the Bayesian approach allows us to both characterize the distribution of staffer effects, and also estimate a baseline model that we will extend in Sections 6-7.

In our baseline specification, we assume that α_i and γ_j are distributed log-normally according to $(\mu_\alpha, \sigma_\alpha)$ and $(\mu_\gamma, \sigma_\gamma)$ respectively. As with all Bayesian estimation, the choice of priors is important. We select weakly informative priors that are designed only to rule out extremely implausible parameter estimates, and rely on the data to inform the posterior. For example, the number of bills co-sponsored by an office in a quarter has a mean of 16.3 and a standard deviation of 18.9. In our baseline specification, the prior for both μ_γ and μ_α is $\mathcal{N}(0, 10)$, meaning that the average staffer producing 0.0001 bills or 10,000 bills are all within 1 standard deviation of the prior mean. We believe that this captures all reasonable levels of potential office and staffer productivity. Priors for $\sigma_\alpha, \sigma_\gamma$ are even more dispersed at $\mathcal{N}(0, 25)$. In Appendix A, we provide a full specification of the priors employed and other estimation details. Throughout the remainder of this paper, we focus our attention on the number of co-sponsored bills produced by an office, given that we are most powered to detect effects for this outcome of interest.

We report the posterior mean, 2.5th, and 97.5th percentiles for the staffer share in Table 3, Panel A.[39] We find that the staffer share of differences is just under 40%, a very similar quantity to that estimated in the dynamic specification (Figure 2, Panel A).[40]

In Table 3, Panel A, and Figure A.8, we present summary statistics and trace plots for the posterior for the distribution of staffer effects and office effects. All signs indicate model convergence. We find posterior means of $(\mu_\alpha, \sigma_\alpha) = (-5.38, 3.38)$ and $(\mu_\gamma, \sigma_\gamma) = (2.52, 0.52)$. This implies that the average staffer produces $\exp(\mu_\alpha + \frac{\sigma_\alpha^2}{2}) = 1.39$ bills per quarter, and that the average office (Representative) produces 14.22 bills per quarter. However, the median staffer produces almost no bills, in line with the fact that the majority of staffers do not work in roles directly tied to legislative work. In Figure 4, we plot the posterior distributions for $\log(\alpha_i)$ and $\log(\gamma_j)$. The density plotted is $\frac{1}{S}\sum_{s=1}^{S} f_{\mu_s,\sigma_s}(x)$ where f is the log-normal density and S is the number of MCMC samples. The histograms plot posterior means for each individual α_i and γ_j. As expected, the vast majority of staffers have virtually no impact on the legislative process. However, there is a meaningful right tail of staffer productivity, including a handful of staffers who have productivities on par with Representatives. In contrast, the distribution of office

[39] These percentiles correspond to a 95% credible set, which under Bernstein-von Mises can also be interpreted as 95% confidence interval.

[40] In Table 3, we also present $\tilde{n}$, the effective sample size of the estimated parameter, and $\hat{R}$, a measure of MCMC convergence (where closer to 1 is better). The effective sample size varies across parameters because it is decreasing in the within-chain correlation of sampled posterior values, which may vary by parameter. For $\hat{R}$, a common rule of thumb is that values beneath 1.05 are acceptable (see `https://mc-stan.org/rstan/reference/Rhat.html`). We present additional properties of the sampler in Table A.5. All diagnostics look regular. We also present a trace plot for the Representative share of the difference in co-sponsored bills between above and below median offices in Figure A.8, Panel A.

Table 3: Bayesian estimates of legislative productivity

Parameter	Mean	2.5%	97.5%	$\tilde{n}$	$\hat{R}$
(1)	(2)	(3)	(4)	(5)	(6)
Panel A: baseline model					
μ_α	-5.38	-7.20	-4.03	680	1.01
μ_γ	2.52	2.43	2.60	647	1.00
σ_α	3.38	2.77	4.18	545	1.01
σ_γ	0.52	0.48	0.57	405	1.00
σ_τ	12.80	9.91	16.42	186	1.03
σ	10.22	10.08	10.38	2361	1.00
Staffer share	0.37	0.34	0.41	1071	1.00
Panel B: nested CES					
ρ_{staff}	0.79	0.62	0.89	4014	1.00
ρ_{rep}	0.13	0.10	0.17	169	1.03
μ_α	-39.96	-51.91	-29.22	1817	1.01
μ_γ	2.36	2.24	2.47	450	1.02
σ_α	9.98	7.50	12.91	1680	1.00
σ_γ	0.54	0.50	0.60	459	1.01
σ_τ	12.65	9.60	16.75	138	1.04
σ	10.05	9.91	10.19	5348	1.00
Staffer share	0.43	0.37	0.47	1133	1.00

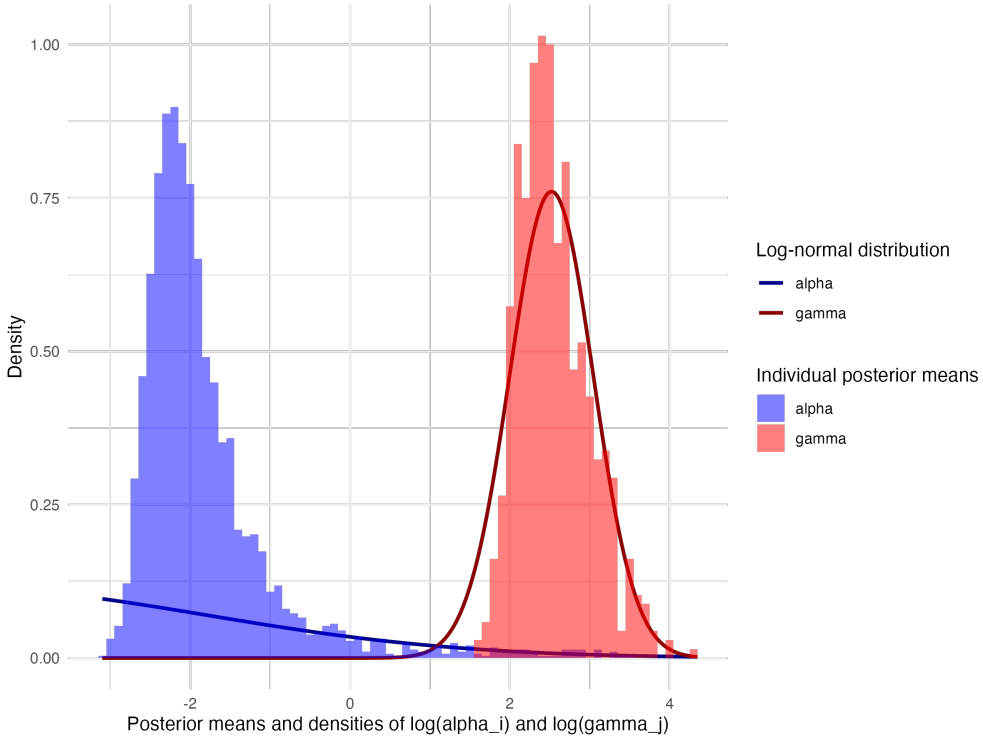
Notes: This table presents parameter estimates from the baseline model (Panel A) and CES model (Panel B). Columns 2 through 4 display posterior means, 2.5th percentiles, and 97.5th percentiles respectively. Column 5 displays the effective sample size and column 6 displays the $\hat{R}$, a measure of MCMC convergence. $\mu_\alpha, \sigma_\alpha$ are the staffer mean and standard deviation, $\mu_\gamma, \sigma_\gamma$ are the office mean and standard deviation. Both parameterize log-normal distributions. σ_τ is the time standard deviation, and σ is the error standard deviation. Rep. share is the share of differences between above and below median bill output offices that is attributable to Representatives. $\rho_{\text{staff}}, \rho_{\text{rep}}$ are elasticities of substitution within a team of staffers, and between the staffers and office, respectively.

productivity is much more concentrated, suggesting that there is much less dispersion in Representative legislative productivity. Differences between the plotted distributions and histograms may indicate when the data has meaningfully informed and overwhelmed any distributional assumptions we have made on α_i, γ_j.

Robustness In Table A.7, Panel A, we replicate our baseline estimate, the staffer share of differences in co-sponsored bills between above and below median offices. In Panel B, we consider alternate mean decompositions. In Panel B.1, we show the staffer share for top versus bottom tercile offices, and in Panel B.2, we show the staffer share for top versus

bottom quartile offices. Examining these alternate differences between offices does not materially change the staffer share.

Figure 4: Posterior distribution for α_i, γ_j from Bayesian estimation



Note: This figure plots the posterior mean of $\log(\alpha_i)$ (in blue) and $\log(\gamma_j)$ (in red) for each staffer and office. The dark lines correspond to the posterior log-normal distribution, whereas the light histograms correspond to posterior means for each individual α_i, γ_j. Estimates for α_i and γ_j come from a Bayesian estimation of Equation 1.

Our baseline estimate assumes a log-normal distribution for both staffer and office effects. In Table A.7, Panel C, we show the staffer share under alternate distributional assumptions. In Panel C.1, we estimate a normal distribution for both staffer and office effects, in Panel C.2, we estimate a Fréchet distribution, and in Panel C.3, we estimate using a mixture of normals instead.[41] All display similar (if not slightly larger) staffer shares to the baseline log-normal distribution.

We also examine the staffer share of outcomes beyond just the number of co-sponsored bills. In Table A.7, Panel D.1 we show the staffer share for co-sponsored bills exiting committee, in Panel D.2 we show the staffer share for co-sponsored bills passing a House vote, in Panel D.3 we show the staffer share for co-sponsored bills becoming law, and in Panel D.4 we show the staffer share for co-sponsored important bills. We find that

[41]The mixture of normals captures the (potentially) bimodal nature of staffer productivities, between those who are involved in legislative work and those who are not. Given the difficulties in simultaneously estimating multiple distributions at once, this model is less likely to have converged.

Bayesian estimation tends to produce larger staffer shares than OLS.

One concern over the entire analysis presented thus far is that it does not account for staffers who never move. For instance, if non-movers have productivities that are negatively correlated with the movers, then it is possible that the staffer share is upwardly biased. We are able to account for these non-movers in our Bayesian analysis, simply by including the non-movers as well when estimating Equation 1. We note that the distribution of effects of staffers who do not move, but do enter or exit the office (i.e. staffers who join or leave Congress during our sample), is identified. For intuition, consider a staffer who leaves an office at time $t + 1$. If we see that the number of bills produced by the office substantially drops from t to $t + 1$, that informs the productivity of the staffer.[42] In Table A.7, Panel E, we show the staffer share when including all non-movers into the estimation as well. As expected, accounting for this missing mass of staffers in the analysis pushes up the staffer share to roughly 65%.[43]

Finally, we note that while methods for conducting a variance decomposition in our team setting have not been developed for the fixed effects estimated via OLS, they can be directly conducted on the random effects from the Bayesian posterior.[44] In Table 4, we decompose the variance at the office level into a staffer component and an office component. We find that staffers are responsible for 37% of the variance, a figure consistent with our other estimates.[45]

Table 4: Variance decomposition

	Share	95% CI
	(1)	(2)
Staffer share	37.04%	[33.11%, 40.87%]
Office share	85.86%	[82.77%, 88.90%]

Notes: Variance decomposition from Bayesian estimates of Equation 1. Column (1) presents staffer and office shares of variance at the office level, and column (2) presents 95% credible intervals.

Taking stock of the linear decomposition, dynamic estimation, and Bayesian estima-

[42]The most substantive assumption underlying this analysis is that staffers who never move, enter, or exit are drawn from a common distribution to the other staffers.

[43]We focus on the movers for the remainder of our analysis, as the much larger number of non-movers compared to movers makes including them computationally expensive.

[44]While Kline et al. (2020) develop a method to correct for the bias created by plug-in methods when estimating variance decompositions in the standard AKM mover design, this has not been extended to the team based mover design.

[45]The staffer and office variances sum to more than 100% due to the correlation between the two. There is also no necessary connection between the shares estimated via variance and mean decompositions.

tion, we find a consistent throughline of staffers accounting for roughly 40% or more of the difference in co-sponsored bills produced between offices. Although estimates for other stages of the bill life-cycle and important bills are often less precise, across these legislative outcomes, we are typically able to bound the staffer share away from 100% (staffers as "puppetmasters") and 0% (staffers as "pawns"). Collectively, these results indicate that staffers play a substantial role in the production of legislation.

6 What Makes Congressional Staffers Effective?

In the analysis above, we have shown that staffers contribute to meaningful differences in legislative productivity across offices. In this section, we dive deeper to understand the mechanisms that determine the efficacy of Congressional staffers. Specifically, we unpack the role that three critical factors might play in determining the impact of staffers: the specific jobs that staffers hold, diversity within offices, and the relationship between the team of staffers and the office (Representative).

6.1 Jobs

Which job titles are most responsible for an office's legislative productivity? The literature typically posits a "job ladder" wherein productivity monotonically increases as one is promoted, justifying increased pay (Moscarini and Postel-Vinay, 2018). We estimate an extension of the model in Equation 1 that allows for jobs to flexibly scale up (or down) the productivity of staffers based on their individual productivity. Specifically, we estimate the model:

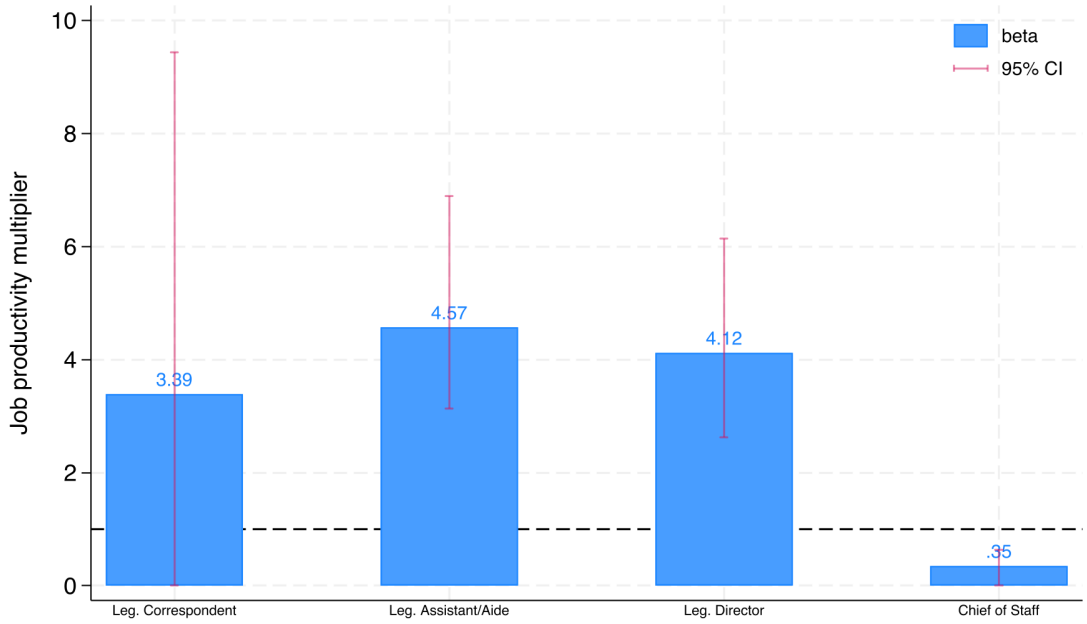
$$y_{jt} = \sum_i \alpha_i \beta_{J(i,t)} m_{ijt} + \gamma_j + \tau_t + \varepsilon_{jt} \tag{9}$$

where $J(i,t)$ refers to the job that a staffer i holds at time t, and β is the new parameter of interest. We separate jobs into four categories based on the legislative staffer career ladder: Legislative Correspondents, Legislative Assistants/Aides, Legislative Directors, and Chiefs of Staff (CoS), as well as an "other" category for all other jobs. We normalize the β for all other jobs to 1.

We present the estimated β parameters in Figure 5. We find that most legislative jobs have a roughly similar amount of productivity (at 3-4 times the typical job), but that there is a sharp drop-off when a staffer is promoted to being a Chief of Staff in their office (down to 0.3). One reason for this may be that Chiefs of Staff are responsible for many different domains, ranging from legislation to communication to regular office administration, and

may not be able to devote as much attention to legislative issues. As such, despite producing legislation being a top priority for many Congressional offices, it appears that the Chief of Staff has much less direct influence over the legislative output of an office when compared to more junior positions.[46] Indeed, the fact that the Chief of Staff coefficient is well below 1 suggests that even holding non-legislative jobs will generally allow a staffer to have greater influence over an office's legislative output than being the Chief of Staff.

Figure 5: Estimates of job-specific β productivities



Note: This figure presents posterior means and 95% credible sets for job-specific slopes on staffer fixed effects (see Equation 9). All other jobs are normalized to have slope 1.

One may also expect that greater experience in Congress also pays returns to individual productivity. In Appendix C.1, we estimate these returns and find that while Representatives gain from their experience (Representatives with 10 full sessions of Congress are about twice as effective as freshmen with no experience), the null of staffer experience being irrelevant cannot be rejected.[47] However, we find that more productive staffers have shorter careers in Congress, supporting concerns that talent may disproportionately leave the Hill for other careers such as lobbying.

[46]The coefficient for LCs (Legislative Correspondents) is imprecisely estimated. This is likely due to the fact that LCs may be given different responsibilities in different offices. Some LCs only handle constituent mail, while others are involved in crafting legislation. Unfortunately, we cannot distinguish different types of LCs from the job title information that we have.

[47]This is a finding consistent with the literature (Crosson et al., 2020).

6.2 Diversity

In April 2024, Congress voted to disband the House Office of Diversity and Inclusion.[48] Proponents and opponents of the measure argued over whether efforts to increase diversity are a "zero-sum game," where the gain of one group must come at the cost of another.[49] Yet diversity has been linked to a host of positive organizational outcomes, from enhanced problem-solving to better alignment with constituent needs (Rock et al., 2016)–especially in complex, high-stakes environments like Congress, where decisions impact millions. To understand how staffer diversity may shape legislative productivity, we examine the association between staffer diversity and legislative productivity, controlling directly for staffer and office quality. To measure diversity, we calculate the share of staffers belonging to each major ethnic group (13 categories, such as "Greater African, Muslim" or "Greater European, British"), gender, educational degree (bachelor's or below, MA, MPP, JD, or other advanced degree), and type of work experience prior to Congress (private sector, policy-facing, local government, or federal government) and compute the Herfindahl-Hirschman Index (HHI) for each office in every quarter along these four measures. We then re-estimate Equation 1, adding each one of these diversity measures in turn as an additional covariate. If diversity is indeed a "zero-sum game," then we would expect to estimate an effect of 0 on each HHI term.

We plot estimated coefficients and 95% credible intervals in Figure 6. For each of these measures, we find that less concentration (more diversity) is associated with more productive offices. Effect sizes are small but statistically significant: moving from the 1st quartile to the 3rd quartile in ethnic diversity increases bills produced by 0.85% of a standard deviation (SD), while the effect is 0.27% of a SD for gender, 0.78% of a SD for education, and 0.28% of a SD for prior work experience. Thus, while we find that offices benefit from diversity, we also rule out extremely large gains from diversity.[50]

6.3 Staffers and Representatives: substitutes or complements?

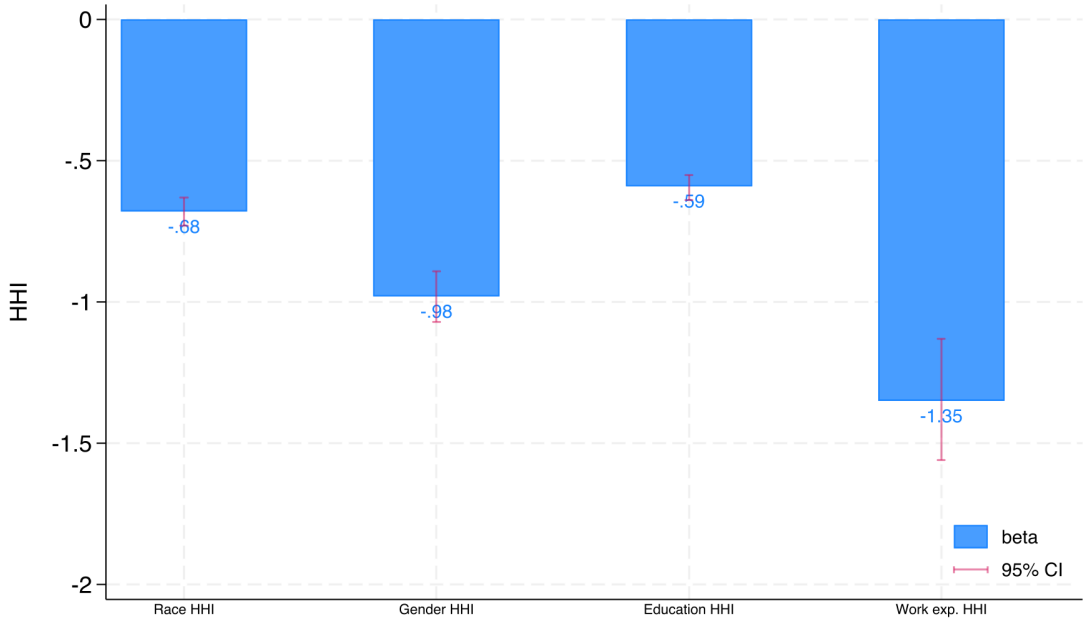
The relationship between staffers and Representatives is at the core of an office's legislative productivity, yet it remains an open question whether they function as complements or substitutes. Does a well-rounded team of skilled staffers enhance the effectiveness of

[48] `https://www.cnn.com/2024/03/22/politics/house-office-diversity-inclusion-disbanded/index.html`

[49] `https://thehill.com/homenews/race-politics/4564129-dei-advocates-sound-alarm-house-office-diversity/`

[50] One reason why this effect may be small could be the allocation of job titles within offices. For instance, Ritchie and You (2021) find that while the gender balance among staffers is roughly equitable, female staffers are much less likely to advance to powerful positions.

Figure 6: Estimates of gains from diversity



Note: This figure presents posterior means and 95% credible sets for the gains to within office HHI (Herfindahl–Hirschman index) for race, gender, education, and prior work experience. Higher HHI corresponds to greater concentration and less diversity.

a strong Representative, or can they fill in gaps for a less experienced one? And within a team, is it better to rely on a single 'superstar' or to distribute responsibilities across several capable staffers?

To answer these questions, we move beyond the assumptions of the linear model imposed in Equation 1, and introduce a more flexible nested CES production function:

$$y_{jt} = \left[\left[\sum_i \alpha_i^{\rho_{\text{staff}}} m_{i,j,t} \right]^{\rho_{\text{office}}/\rho_{\text{staff}}} + \gamma_j^{\rho_{\text{office}}} \right]^{1/\rho_{\text{office}}} + \tau_t + \varepsilon_{jt} \tag{10}$$

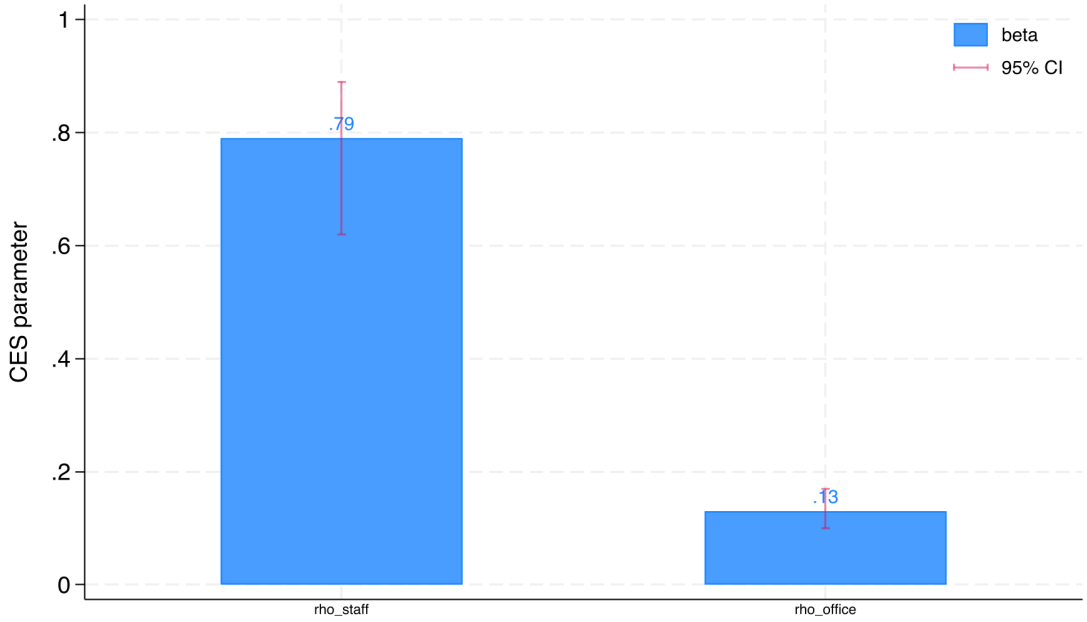
where ρ_{staff} and ρ_{office} are the parameters of interest. Specifically, ρ_{staff} governs the degree of substitutability of staffers within a team, while ρ_{office} governs the degree of substitutability between the team of staffers and the office (or Representative).

We present estimated coefficients and 95% credible intervals in Figure 7.[51] Recall that under perfect substitutes $\rho = 1$, and that under Cobb-Douglas $\rho = 0$. We find that staffers within a team are highly substitutable for one another. The posterior mean for

[51] Table 3, Panel B contains posterior means, 2.5th and 97.5th percentiles for various parameters of interest. Trace plots can be found in Figure A.9. We note that the estimated staffer share under this more flexible production function is similar to prior estimates. The same is true for other models estimated in this section (results available upon request).

$\rho_{\text{staff}} = 0.79$.[52] This suggests that maximizing the sum of staffer productivities α_i is close to optimal, regardless of whether the skill is highly concentrated within one individual or spread over many. On the other hand, we find that the production function for the team of staffers and the office is close to Cobb-Douglas. The posterior mean for $\rho_{\text{office}} = 0.13$.[53] Heuristically speaking, this means that an office consisting of a superstar Representative but very unproductive staffers and an office consisting of a very unproductive Representative and superstar staffers will both be less productive than an office containing a more balanced mix of skill between Representatives and staffers. Thus, Representatives facing a hiring trade-off between staffer compensation and staffer skill should hire staffers that tend to reflect their own productivity.

Figure 7: Estimates of nested CES ρ parameters



Note: This figure presents posterior means and 95% credible sets for nested CES ρ parameters (see Equation 10). ρ_{staff} governs the degree of substitutability between staffers on a team, while ρ_{office} governs the degree of substitutability between the team of staffers and the office (Representative).

Identifying Pareto improvements Do Representatives generally follow this guideline? We formalize this by testing for potential Pareto improvements, circumstances under which two offices may trade a staffer and both are left better off after the trade. We operationalize this by additionally assuming that offices have access to a technology that produces bills with a linear cost function. One interpretation of this technology could

[52] This translates to $\sigma = \frac{1}{1-\rho} = 4.76$.

[53] This translates to $\sigma = 1.15$.

be Congressional interns—offices can hire interns for variable lengths of time (and pay wages proportional to time worked) to assist with the production of legislation. This generates the price of a bill, p.[54]

Consider two offices A, B which employ the set of staffers $\mathcal{A}$ and $\mathcal{B}$ respectively. Let w_a denote the wage of staffer a. Then a Pareto improvement between the offices is possible if there exists $a \in \mathcal{A}$ and $b \in \mathcal{B}$ such that the following holds:

$$\left[\alpha_a^{\rho_{\text{staff}}} + \sum_{a' \in \mathcal{A}, a' \neq a} \alpha_{a'}^{\rho_{\text{staff}}}]^{\rho_{\text{office}}/\rho_{\text{staff}}} + \gamma_A^{\rho_{\text{office}}}\right]^{1/\rho_{\text{office}}} - \frac{w_a}{p} \leq [[\alpha_b^{\rho_{\text{staff}}} + \sum_{a' \in \mathcal{A}, a' \neq a} \alpha_{a'}^{\rho_{\text{staff}}}]^{\rho_{\text{office}}/\rho_{\text{staff}}} + \gamma_A^{\rho_{\text{office}}}]^{1/\rho_{\text{office}}} - \frac{w_b}{p}$$

$$\left[\alpha_b^{\rho_{\text{staff}}} + \sum_{b' \in \mathcal{B}, b' \neq b} \alpha_{b'}^{\rho_{\text{staff}}}]^{\rho_{\text{office}}/\rho_{\text{staff}}} + \gamma_B^{\rho_{\text{office}}}\right]^{1/\rho_{\text{office}}} - \frac{w_b}{p} \leq [[\alpha_a^{\rho_{\text{staff}}} + \sum_{b' \in \mathcal{B}, b' \neq b} \alpha_{b'}^{\rho_{\text{staff}}}]^{\rho_{\text{office}}/\rho_{\text{staff}}} + \gamma_B^{\rho_{\text{office}}}]^{1/\rho_{\text{office}}} - \frac{w_a}{p}$$

with at least one inequality strict. We test for all possible Pareto improving trades under the assumption that $p = 50,000$ ($50,000 of intern time produces an additional bill), restricting attention to only trades of legislative staff between offices.

Out of roughly 10 million potential trades of staffers between offices, we find that there are only 897 trades that lead to Pareto improvements. We conduct a permutation test, randomly shuffling a Representative's assigned office 100 times and computing the number of Pareto improving trades. We find that the empirically estimated number of improvements is smaller than every value from the permutation test. The small number of potential improvements suggests that offices hire (legislative) staffers with some degree of sophistication. Both the team's productivity and the office's underlying productivity must be accurately incorporated into a hiring decision for this result to hold. The relatively small number of Pareto improvements also suggests that the outcome we examine, bills, is not an unreasonable proxy for an office's (legislative) objective function.[55]

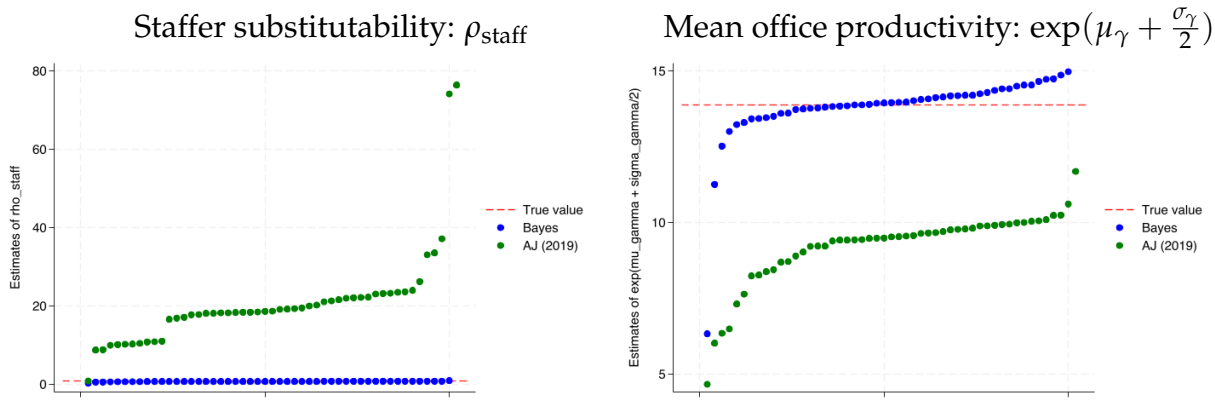
A comparison of methods Although most existing methods for estimating individual contributions to team output are not feasible in our context (see Appendix B), Ahmadpoor and Jones (2019) provide a feasible method. Their estimation routine uses gradient descent to solve a nonlinear least squares problem, namely recovering individual effects in a CES production function. To compare their method against our Bayesian procedure, we first generate 50 simulated datasets that use the same history of moves as our primary sample, and simulate staffer and Representative effects using the same distributions from our baseline estimates in Table 3, Panel B. We then compare estimates from both procedures. In Figure 8, we plot estimated values of staffer substitutability (ρ_{staff}) and mean of-

[54] Unlike staffers, intern skill is typically ex-ante unobserved. Assuming that intern skill is drawn from a common distribution across offices thus supports a common price for bills p, unlike staffer wages which may individually vary.

[55] To the best of our knowledge, this is also a novel test to validate AKM-style models.

fice productivity ($\exp(\mu_\gamma + \frac{\sigma_\gamma}{2})$), alongside the true parameter value. While the Bayesian method is randomly initialized, the nonlinear least squares method is initialized with the Oracle solution. Despite this asymmetry, we find that the Bayesian method is better at estimating these parameters, showing that there are instances where our approach may be preferred.

Figure 8: Comparing estimates of a nested CES model across methods



Note: This figure plots estimates of staffer substitutability (ρ_{staff}) and mean office productivity ($\exp(\mu_\gamma + \frac{\sigma_\gamma}{2})$) from a nested CES model. Data are generated 50 times from distributions parameterized by the posterior means in Table 3. Parameters are estimated using a random effects Bayesian method, presented in blue, and the nonlinear least squares method of Ahmadpoor and Jones (2019), in green. The Bayesian method is randomly initialized, while the nonlinear least squares method is initialized with the Oracle solution. The true value is plotted as a dashed red line. Estimates are ordered by size.

7 Estimating Staffer Ideology

While we have demonstrated that staffers play a crucial role in legislative productivity, an equally important question is whether they influence the ideological content of bills. Staffers are not mere administrative support; they bring their own viewpoints and values, which may shape legislative content. In this section, we estimate the ideology of Representatives and staffers through the content of the bills they write. We conjecture that the ideological content of the bills produced by a Congressional office is the weighted average of all individual's ideological ideal points within the office:

$$\iota_{jt} = \frac{\sum_i \omega_i \iota_i m_{i,j,t} + \omega_j \iota_j}{\sum_i \omega_i m_{i,j,t} + \omega_j} + \varepsilon_{jt} \tag{11}$$

where ω is a weight and ι is the staffer or office's ideal point. To estimate the equation above, we take existing estimates of bill ideology from DW-Nominate and compute the

average ideology for bills co-sponsored by an office in a given quarter. This is the outcome ι_{jt}. We take our estimates of the staffer and office effects under the baseline Bayesian model as weights, so that $\omega_i = \alpha_i, \omega_j = \gamma_j \forall i, j$—the assumption being that the producer of a bill can control the ideological content of the bill that they produce—and estimate the model using Hamiltonian Monte Carlo. This produces estimates of ideology ι_i and ι_j. The distinction between our estimates and DW-Nominate is that while Nominate measures ideology during roll call voting, our estimates measure ideology when writing bills. Alongside Representatives, we are additionally able to estimate the ideology of staffers.

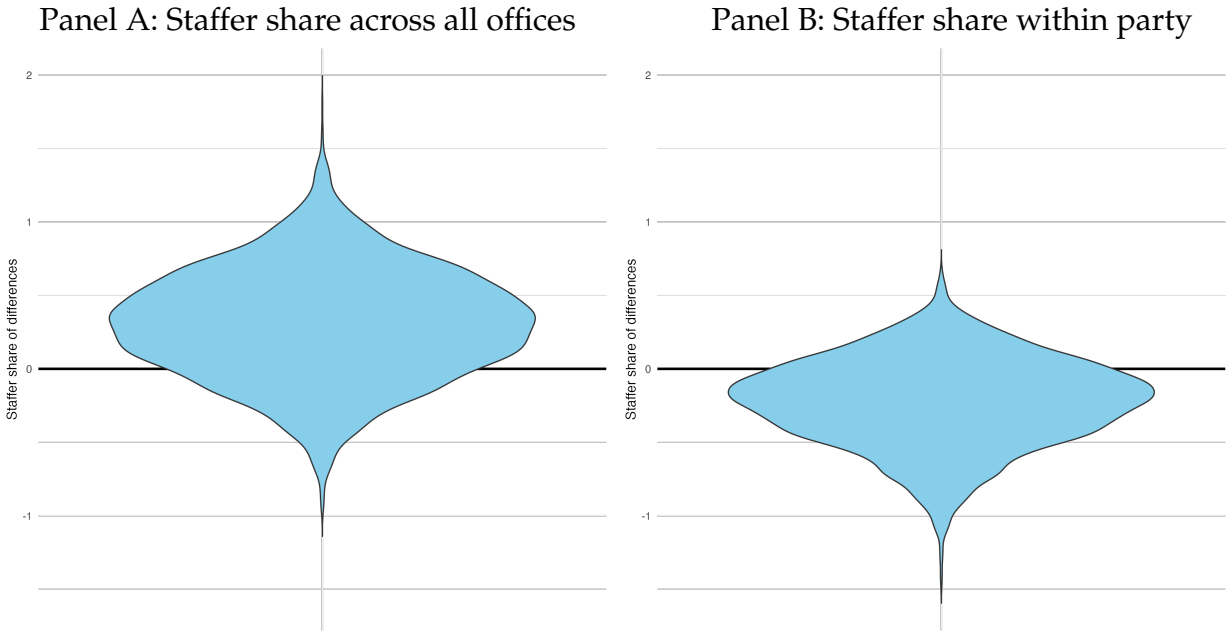
We present estimates of the staffer share of differences between above and median offices along the first dimension of DW-Nominate in Figure 9.[56] The first dimension of DW-Nominate corresponds to the typical liberal/conservative or left/right ideological divide in the US. We find that Representatives explain 70% of the differences between offices along this first dimension, suggesting that they are more responsible for polarization than staffers. This finding grows more stark once we look within parties: approximately all of the within-party differences in ideology along the first DW-Nominate dimension are due to the Representative. Staffers have little influence along typical partisan lines. If anything, the posterior mean for the staffer share being -22% suggests that staffers play a moderating influence on their Representatives; in their absence, the bills produced by Congress would be 22% more extreme. To illustrate, we estimate that Representative Yvette Clark (D-NY) writes some of the most ideologically liberal bills (in the 98th percentile), while her staff is relatively moderate (in the 58th percentile among Democrat staff). Similarly, while Representative Ralph Norman (R-SC) writes some of the most ideologically conservative bills (in the 99th percentile), his staff is also relatively moderate (in the 57th percentile among Republican staff). In Figure A.11, we plot posterior means of Representative ideology and staffer ideology. Staffer and Representative ideologies are not clearly correlated within party.[57] Thus, we suggest that Representatives at ideological extremes of their parties may be reigned in by their staffs.

What could explain the moderating impact that staffers have on partisan politics? One possibility is that staffers work in teams: while individual staffers might be extreme, the combination of such individuals in a team could balance out, resulting in a neutral overall effect. However, this does not appear to be the case. The standard deviation of staffer ideology is 0.184, whereas the standard deviation of staffer ideologies at the Representa-

[56] Trace plots are presented in Figure A.10.

[57] One may be concerned that the estimates of bill ideology from DW-Nominate reflect cut-points and not bill proposal locations. To address this, we estimate bill proposal locations following Woon (2008) and estimate a posterior mean for the staffer share of 1.5% across parties and -3.8% within party, with 95% confidence sets [-3.9%,7.0%] and [-9.1%, 1.9%] respectively. For details, see Appendix B.4.

Figure 9: Violin plots for the staffer share of differences in ideology



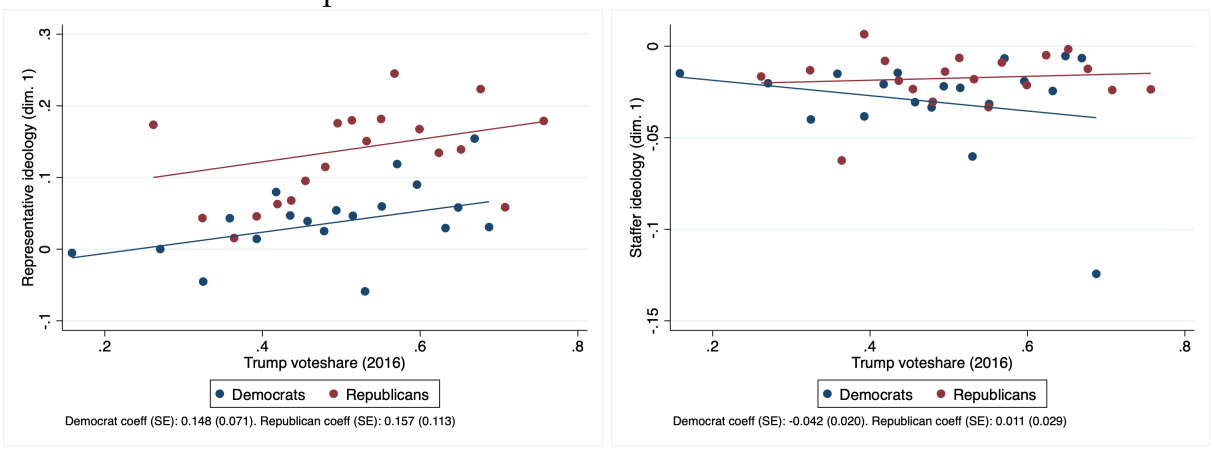
Note: This figure presents violin plots for a Bayesian estimation of the ideology of bills produced by an office (Equation 11), which is a weighted average of office (Representative) and staffer ideal points. Weights are estimated from the team-based mover design (Equation 1). Both panels show the staffer share of differences between above and below-median offices in ideology for the first DW-Nominate dimension, the conventional left-right partisan ideology. Panel A shows the staffer share of differences across all offices, while Panel B shows the staffer share of differences for above and below-median offices within a party.

tive level (using a productivity weighted average) is 0.187. In other words, aggregating staffers into teams does not reduce the overall variance in ideology, implying that other factors are responsible for the moderating influence we observe.

Our favored explanation centers on the distinct institutional roles and incentives of staffers versus Representatives. Unlike Representatives, staffers do not directly face electoral incentives. This insulation from electoral accountability may allow them to prioritize policy expertise over strict ideological adherence. We hypothesize that staffers have a moderating influence on partisanship because of this. In Figure 10, we plot binscatters of the voteshare for Trump (2016) in a Representative's district against their ideology, for both Democrats and Republicans. In Panel A, we show that the Trump voteshare strongly predicts the ideology of Representatives, with a similar slope across parties. This strong correlation suggests that Representatives align their ideological positions closely with the partisan preferences of their constituencies, reflecting responsiveness to voter demand. In Panel B, we find that the Trump voteshare is much less predictive of staffer ideology. Among Republican staffers, the correlation is not statistically significant, while for Democrat staffers, the relationship reverses slightly: staffers in districts that leaned Trump are

slightly more liberal. This divergence aligns with a theory of electoral incentives, but is also compatible with a story where Representatives are selected from the population they represent (unlike staffers, who may be recruited from different districts). To separate these stories, in Figure A.12, we regress Representative and staffer ideology against the Representative's own voteshare, as a proxy for the "safeness" of a Representative's seat. For both Democrats and Republicans, we find that Representatives in safer seats become more ideologically extreme. However, this pattern is less pronounced and more ambiguous for staffers. Across both parties, we find that the slope on Representatives is steeper than that for staffers, providing further evidence that Representatives are more elastic to these electoral concerns. Thus, while not wholly conclusive, our preferred interpretation of these results is that polarization among Representatives is primarily driven by voter demand, whereas staffers, shielded from electoral incentives, may act as a moderating force in the legislative process.

Figure 10: Representative and staffer ideologies by Trump voteshare



Note: This figure presents binscatter plots for posterior means of Representative (left) and staffer (right) estimated first dimension ideologies against their district's Trump voteshare in 2016. Positive values are more conservative. Regressions are at the Representative level, with staffer ideologies the (productivity-weighted) average across staffers who work for the Representative. Democrat are plotted in blue, Republicans in red. Standard errors for the regression coefficients presented are robust.

Next, we turn to the second dimension of DW-Nominate. The staffer share of differences are plotted in Figure A.15. In contrast to the first dimension, we find that staffers account for the lion's share of differences in ideology along the second dimension. The posterior mean for this staffer share is 78% across all offices and 84% within party. Though the second dimension has proven difficult to interpret by scholars, it is orthogonal to the first dimension's left/right ideology by construction. We find that staffers exert more influence when focusing on these alternate, less partisan issues. This is consistent with Weber's conception of bureaucracy (Weber, 1968), or Hanson's notion of pulling policy

ropes sideways: an effective bureaucrat that focuses on policy over politics can plausibly make a bigger impact, as they sidestep a frequently gridlocked partisan debate.[58] Finally, to verify the results of our estimation process, we regress the posterior means of a Representative's estimated ideology when writing bills against their DW-Nominate score from roll call votes. Table A.6 shows that these measures tend to be significantly correlated.

Taken together, these results suggest that staffers are most impactful when they pursue policy goals outside of the typical partisan debate, and that because they are not as beholden to electoral incentives, they may act as a moderating force in partisan politics.

8 Conclusion

In this paper, we have laid out a new framework for estimating the impact of individuals working within teams when individual contributions are unobserved. Through linear decompositions, dynamic estimates, and Bayesian methods, we have quantified the crucial role that Congressional staffers play in the legislative process. Our findings indicate that staffers are not merely extensions of their members of Congress but are influential contributors to legislative productivity, accounting for 40% or more of the variation in productivity between different Congressional offices. In extensions to our baseline model, we further characterize the importance of staffers, identifying the legislative positions in which they have the greatest impact, the complementarity between staffer and Representative productivity, and the ways in which they can moderate ideological extremism and contribute to less partisan legislation. Collectively, our results suggest ways in which staffers new to Congress can quickly make an impact: by taking on legislative roles, working for Representatives with high intrinsic productivity, and focusing on policy issues outside of conventional partisan divides.

Looking forward, there are many interesting questions to explore. First, the influence of staffers likely reaches beyond legislative productivity to other areas such as constituent services, media communication, and government oversight, as hinted in decompositions of Congressional office spending (see Figure A.2). What is the impact of Congressional staffers in these domains? Future research may investigate this question. Second, although this paper has focused on the impact of personal staff, committee staff also play a large role in the production of legislation. It would be interesting to understand the role that other staffers, and other non-elected members of government more broadly, play in the legislative process. Third, this paper raises the question of whether staffers are able to reduce polarization and make progress on non-partisan issues precisely because they

[58] See `https://www.overcomingbias.com/p/policy_tugowarhtml`.

are non-elected. This insulation from electoral accountability may enable them to act as mediators, facilitating bipartisan cooperation and progress on less partisan issues. The role of electoral accountability on political polarization more broadly speaking could be an important avenue for further research.

To conclude, we present one final speculative regression on the role that Congressional staffers play in our political system. Representatives in the House were first allowed to hire a single staffer in 1893. Over the course of the next century, Congress expanded the maximum number of staffers that Representatives could hire fifteen times. In Figure A.13, we present coefficients from a two-way fixed effect event study, regressing the maximum number of staffers that Representatives can hire on the probability that the incumbent wins an election. We find that a doubling in the number of staffers corresponds to a cumulative roughly 20% increase in the probability that an incumbent wins. It appears that staffers even influence the thing that politicians may care about most: re-election.

References

Aberbach, Joel D, Robert D Putnam, and Bert A Rockman, *Bureaucrats and politicians in western democracies*, Harvard university press, 1981.

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

_ , _ , and Simon Woodcock, "Econometric analyses of linked employer–employee data," in "The econometrics of panel data: Fundamentals and recent developments in theory and practice," Springer, 2008, pp. 727–760.

_ , Kevin L McKinney, and Ian M Schmutte, "Modeling endogenous mobility in wage determiniation," *US Census Bureau Center for Economic Studies Paper No. CES-WP-15-18*, 2015.

Ahmadpoor, Mohammad and Benjamin F Jones, "Decoding team and individual impact in science and invention," *Proceedings of the National Academy of Sciences*, 2019, *116* (28), 13885–13890.

Andrews, Martyn J, Len Gill, Thorsten Schank, and Richard Upward, "High wage workers and low wage firms: negative assortative matching or limited mobility bias?," *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2008, *171* (3), 673–697.

Aneja, Abhay and Guo Xu, "Strengthening State Capacity: Civil Service Reform and Public Sector Performance during the Gilded Age," *The American Economic Review*, Forthcoming.

Angrist, Joshua D, Peter D Hull, Parag A Pathak, and Christopher R Walters, "Leveraging lotteries for school value-added: Testing and estimation," *The Quarterly Journal of Economics*, 2017, *132* (2), 871–919.

Arora, Abhishek and Melissa Dell, "LinkTransformer: A Unified Package for Record Linkage with Transformer Language Models," *arXiv preprint arXiv:2309.00789*, 2023.

Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter, "Management practices, workforce selection, and productivity," *Journal of Labor Economics*, 2018, *36* (S1), S371–S409.

Bergeron, Augustin, Pedro Bessone, John Kabeya Kabeya, Gabriel Z Tourek, and Jonathan L Weigel, "Optimal assignment of bureaucrats: Evidence from randomly assigned tax collectors in the DRC," Technical Report, National Bureau of Economic Research 2022.

Bertrand, Marianne, Matilde Bombardini, and Francesco Trebbi, "Is It Whom You Know or What You Know? An Empirical Assessment of the Lobbying Process," *American Economic Review*, December 2014, *104* (12), 3885–3920.

Besley, Timothy and Torsten Persson, "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, August 2009, *99* (4), 1218–1244.

Bonhomme, Stéphane, "Teams: Heterogeneity, sorting, and complementarity," *Proceedings of the 15th World Congress of the Econometric Society*, 2021.

_ , Thibaut Lamadon, and Elena Manresa, "A distributional framework for matched employer employee data," *Econometrica*, 2019, *87* (3), 699–739.

Boxell, Levi, Matthew Gentzkow, and Jesse M Shapiro, "Cross-country trends in affective polarization," *Review of Economics and Statistics*, 2024, *106* (2), 557–565.

Cantoni, Enrico and Vincent Pons, "Does context outweigh individual characteristics in driving voting behavior? Evidence from relocations within the United States," *American Economic Review*, 2022, *112* (4), 1226–1272.

Card, David, Jörg Heining, and Patrick Kline, "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly journal of economics*, 2013, *128* (3), 967–1015.

Chan, David C, "Influence and information in team decisions: evidence from medical residency," *American Economic Journal: Economic Policy*, 2021, *13* (1), 106–137.

Chen, Yiqun, "Team-specific human capital and team performance: evidence from doctors," *American economic review*, 2021, *111* (12), 3923–3962.

Chetty, Raj and Nathaniel Hendren, "The impacts of neighborhoods on intergenerational mobility II: County-level estimates," *The Quarterly Journal of Economics*, 2018, *133* (3), 1163–1228.

_ , John N Friedman, and Jonah E Rockoff, "Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates," *American economic review*, 2014, *104* (9), 2593–2632.

Childers, David, Jesús Fernández-Villaverde, Jesse Perla, Christopher Rackauckas, and Peifan Wu, "Differentiable state-space models and hamiltonian monte carlo estimation," Technical Report, National Bureau of Economic Research 2022.

Constantine, Noah and Sergio Correia, "Putting the "i" in Fixed Effects: Individual Fixed Effects with Group Level Outcomes," Technical Report, Working paper 2021.

Cross, Jesse M and Abbe R Gluck, "The Congressional Bureaucracy," *University of Pennsylvania Law Review*, 2020, *168* (6), 1541–1683.

Crosson, Jesse M, Alexander C Furnas, Timothy Lapira, and Casey Burgat, "Partisan competition and the decline in legislative capacity among congressional offices," *Legislative Studies Quarterly*, 2021, *46* (3), 745–789.

_ , Geoffrey M Lorenz, Craig Volden, and Alan E Wiseman, "How experienced legislative staff contribute to effective lawmaking," in "Congress overwhelmed: The decline in congressional capacity and prospects for reform" University of Chicago Press Chicago, IL 2020, pp. 209–24.

Dahlström, Carl and Victor Lapuente, "Comparative bureaucratic politics," *Annual Review of Political Science*, 2022, *25*, 43–63.

DeGregorio, Christine, "Professionals in the US Congress: an analysis of working styles," *Legislative Studies Quarterly*, 1988, pp. 459–476.

Dikkala, Nishanth, Greg Lewis, Lester Mackey, and Vasilis Syrgkanis, "Minimax estimation of conditional moment models," *Advances in Neural Information Processing Systems*, 2020, *33*, 12248–12262.

Enders, Adam M, "Issues versus affect: How do elite and mass polarization compare?," *The Journal of Politics*, 2021, *83* (4), 1872–1877.

Ferrara, Andreas, Patrick A Testa, and Liyang Zhou, "New Area-and Population-Based Geographic Crosswalks for US Counties and Congressional Districts, 1790–2020," Technical Report, National Bureau of Economic Research 2024.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Place-based drivers of mortality: Evidence from migration," *American Economic Review*, 2021, *111* (8), 2697–2735.

Flynn, Daniel J and Laurel Harbridge, "How partisan conflict in congress affects public opinion: Strategies, outcomes, and issue differences," *American Politics Research*, 2016, *44* (5), 875–902.

Fox, Harrison W and Susan Webb Hammond, *Congressional staffs: The invisible force in American lawmaking*, New York: The Free Press, 1977.

Freund, Lukas, "Superstar Teams," *Available at SSRN 4312245*, 2024.

Gailmard, Sean and John W Patty, "Formal models of bureaucracy," *Annual Review of Political Science*, 2012, *15*, 353–377.

Gelman, Andrew, John B Carlin, Hal S Stern, and Donald B Rubin, *Bayesian data analysis*, Chapman and Hall/CRC, 1995.

Hagedorn, Sara L, "Taking the lead congressional staffers and their role in the policy process." PhD dissertation, University of Colorado at Boulder 2015.

Hall, Richard L and Alan V Deardorff, "Lobbying as legislative subsidy," *American Political Science Review*, 2006, *100* (1), 69–84.

Hoffman, Mitchell and Christopher T Stanton, "People, Practices, and Productivity: A Review of New Advances in Personnel Economics," Working Paper 32849, National Bureau of Economic Research August 2024.

i Vidal, Jordi Blanes, Mirko Draca, and Christian Fons-Rosen, "Revolving Door Lobbyists," *American Economic Review*, December 2012, *102* (7), 3731–48.

Iranzo, Susana, Fabiano Schivardi, and Elisa Tosetti, "Skill dispersion and firm productivity: An analysis with employer-employee matched data," *Journal of Labor Economics*, 2008, *26* (2), 247–285.

Kingdon, John W, *Congressmen's voting decisions*, University of Michigan Press, 1989.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten, "Leave-out estimation of variance components," *Econometrica*, 2020, *88* (5), 1859–1898.

Lazarus, Jeffrey, Joshua McCrain, and Amy McKay, "Lobbying access, congressional employment, and the determinants of revolving door lobbying," 2019.

Leal, David L and Frederick M Hess, "Who chooses experience? Examining the use of veteran staff by house freshmen," *Polity*, 2004, *36* (4), 651–664.

Mayhew, David R, *Congress: The Electoral Connection*, Yale university press, 2004.

Montgomery, Jacob M and Brendan Nyhan, "The effects of congressional staff networks in the US House of Representatives," *The Journal of Politics*, 2017, *79* (3), 745–761.

Moscarini, Giuseppe and Fabien Postel-Vinay, "The Cyclical Job Ladder," *Annual Review of Economics*, 2018, *10* (1), 165–188.

Neyman, Jerzy and Elizabeth L Scott, "Consistent estimates based on partially consistent observations," *Econometrica*, 1948, pp. 1–32.

Niskanen, William A, "The peculiar economics of bureaucracy," *The American Economic Review*, 1968, *58* (2), 293–305.

Ommundsen, Emily Cottle, "The Institution's Knowledge: Congressional Staff Experience and Committee Productivity," *Legislative Studies Quarterly*, 2023, *48* (2), 273–303.

Ritchie, Melinda N and Hye Young You, "Women's advancement in politics: Evidence from congressional staff," *The Journal of Politics*, 2021, *83* (2), 421–438.

Rock, David, Heidi Grant, and Jacqui Grey, "Diverse teams feel less comfortable—and that's why they perform better," *Harvard Business Review*, 2016, *95* (9), 22.

Romzek, Barbara S. and Jennifer A. Utter, "Congressional Legislative Staff: Political Professionals or Clerks?," *American Journal of Political Science*, 1997, *41* (4), 1251–1279.

Salisbury, Robert H. and Kenneth A. Shepsle, "Congressional Staff Turnover and the Ties-That-Bind," *American Political Science Review*, 1981, *75* (2), 381–396.

Shepherd, Michael E and Hye Young You, "Exit strategy: career concerns and revolving doors in Congress," *American Political Science Review*, 2020, *114* (1), 270–284.

Silver, David, "Haste or waste? Peer pressure and productivity in the emergency department," *The Review of Economic Studies*, 2021, *88* (3), 1385–1417.

Smith, Steven S, Jason M Roberts, and Ryan J Vander Wielen, *The American Congress*, Cambridge University Press, 2013.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter, "Firming up inequality," *The Quarterly journal of economics*, 2019, *134* (1), 1–50.

Spenkuch, Jörg, Edoardo Teso, and Guo Xu, "Ideology and Performance in Public Organizations," *Econometrica*, 2023.

Vogler, Jan P, "Bureaucracies in historical political economy," in "The Oxford Handbook of Historical Political Economy," Oxford University Press, 2022, pp. C19–S1.

Volden, Craig and Alan E Wiseman, "Legislative effectiveness in Congress," *Manuscript, The Ohio State University*, 2009.

Weber, Max, *Economy and society*, New York: Bedminster Press, 1968.

Weidmann, Ben, Joseph Vecci, Farah Said, David J Deming, and Sonia R Bhalotra, "How Do You Find a Good Manager?," Technical Report, National Bureau of Economic Research 2024.

Wielen, Ryan J Vander and Michael J Vander Wielen, "Unpacking the unknown: a method for identifying status quo distributions," *Public Choice*, 2020, *182* (1), 49–72.

Woon, Jonathan, "Bill sponsorship in Congress: the moderating effect of agenda positions on legislative proposals," *The Journal of Politics*, 2008, *70* (1), 201–216.

Online Appendix for: Puppetmasters or Pawns? The Power of Congressional Staffers

This appendix contains additional material, including figures and tables, for the article "Puppetmasters or Pawns? The Power of Congressional Staffers."

Appendix A Data

Constructing a panel of staffers We construct panel data for staffers, associating each staffer with a single office in each quarter, using the following procedure:

1. We drop all staffers that never move.
2. We drop all staffers who are interns or part-time workers.
3. We keep only staffers that stay in Congress for at least one year (4 quarters).
4. We drop any staffers that work for more than 5 offices in a single quarter. These tend to be staffers in generalist positions (such as technology consultants) who are outside the main focus of our study.
5. For each staffer that works for multiple offices in a quarter, we assign an office to a staffer based on the longest duration worked. We break ties randomly.
6. If a staffer has spent less than 40% of a quarter working for the office, we also drop them.
7. We identify a job as a continuous period of time spent working for the same office, with up to 1 quarter of interruption in the middle.

Important bills and the LES We quote the Center for Effective Lawmaking on their procedure for constructing important bills (which they term "substantive and significant"):

> A bill is deemed substantive and significant if it was introduced during the 93rd – 113th Congresses, and had been the subject of an end-of-the-year write-up in the Congressional Quarterly Almanac. For bills that were introduced into the 114th to 116th Congresses, a bill is deemed substantive and significant if it was mentioned on two or more occasions in the stories published in Congressional Quarterly Weekly/CQ Magazine during that Congress. For bills that were introduced into the 117th Congress and subsequent congresses, a bill

> is deemed substantive and significant if it was mentioned on two or more occasions in the stories published in Congressional Quarterly Weekly/CQ Magazine during that Congress, identified as the subject of a key vote by Project Vote Smart, and/or discussed in the Fact Sheets in the CQ House section of CQ.com.

Many important bills have large numbers of co-sponsors (often in the dozens). Given that we have no effective way to determine which offices were primarily responsible for crafting the original text of these bills, we drop all important bills that received five or more original co-sponsors.

The Legislative Effectiveness Score (LES, Volden and Wiseman (2009)) is constructed as a weighted average of bills introduced, bills receiving committee action, bills passing the house, and bills becoming law, weighted by their importance. The benchmark is the predicted LES from an OLS regression containing covariates Representative seniority and dummies for party majority, and committee and subcommittee chair.[1]

Measures of diversity To assess diversity at the office level, we develop a set of diversity measures across race, gender, educational background, and work experience using the Herfindahl-Hirschman Index (HHI). We begin by gathering staffer-level data from LegiStorm, a comprehensive database of congressional staffers. We use staffer names from the Congressional Statements of Disbursement to search for corresponding profiles on LegiStorm. Staffers who have profiles on LegiStorm are considered successfully matched with the names from the Congressional Statements of Disbursement. LegiStorm provides staffer information including gender, work history, and degrees attained, all of which can be systematically scraped for each individual. This allows us to construct detailed profiles for our analysis.

We construct data on race by employing an open-source tool called ethnicolr[2], which predicts race and ethnicity based on staffer names. This tool is trained on census names and is useful when race information is not available.

To quantify educational diversity, we parse the educational backgrounds listed on LegiStorm profiles. By identifying and categorizing keywords such as "BS", "BA", "Bachelor", and equivalent terms, we classify the degrees attained by each staffer. Degrees are categorized into levels such as bachelor's, master's, MBA, MD, JD, and PhD. This classification enables us to gauge the range and concentration of educational achievements within an office.

In terms of work experience, we leverage the capabilities of GPT-4 to automate and refine the categorization process. We first define categories for work experiences, including internship, federal government, state government, policy, and private sector roles. We then input the entire work history of staffers, focusing on experiences prior to their first congressional role (excluding congressional internships), and use GPT-4 to classify these experiences according to the predefined categories. While language models can sometimes generate unreliable information, this specific classification task is well-suited to GPT-4's capabilities since it involves categorizing existing structured data rather than

[1]The specific algorithm used is outlined in `https://thelawmakers.org/methodology`.

[2]More details of this package can be found at `https://github.com/appeler/ethnicolr`.

generating new content. We also implement a human review process to validate the classifications and ensure accuracy.

Finally, we compute the HHI for each of the measures described above at the office-quarter level.

Election data For Representatives' own electoral performance, we use district-level returns from U.S. House elections from the MIT Election Lab. We calculate vote shares as the proportion of total votes received by each candidate and match these directly to Representatives based on name and district-year identifiers.

For Trump's 2016 district-level performance (Trump voteshare), we start with county-level presidential election returns from the MIT Election Lab. Since these data are at the county level but we need district-level measures, we use the population-based crosswalks developed by Ferrara et al. (2024) to convert county-level vote totals into congressional district-level aggregates. We then calculate Trump's vote share at the district level and match this 2016 baseline measure to the corresponding district each Representative serves.

Appendix B Estimation

Appendix B.1 Linear decomposition

Main model Under the typical mover design model, the number of fixed effects is constant per observation (oftentimes three—one for the individual, one for the group, and one for time). Common fixed effect estimation routines are able to quickly compute the full set of fixed effects through a variety of optimization techniques. However, under the team-based mover design model, the number of fixed effects to estimate varies per observation: a team with eleven staffers has five more fixed effects than a team with six staffers. Although recent progress has been made in estimation when these fixed effects are incidental parameters (Constantine and Correia, 2021), the recovery of the actual fixed effects in a computationally efficient way is an open problem. As a result, to estimate the fixed effects in Equation 1, we simply use the full design matrix containing indicators for all staffers and offices when estimating this model, applying Frisch-Waugh Lowell to residualize any time fixed effects and other incidental parameters.

After applying the standard OLS matrix algebra, we compute the Representative share of differences between high and low legislative output offices and subtract it from 1 to obtain the staffer share. We use the Bayesian bootstrap to compute confidence intervals, weighting each observation by the square root of a random variable distributed $\exp(1)$. Throughout the paper, we use 100 iterations of the bootstrap.

Instrumental variables Our instrumental variable strategy uses two sets of instruments for staffer moves at time t:

1. Turnover: This is an indicator equal to 1 if the origin office lost their election at $t-1$ and the destination office is in the same state as the origin office.

2. Migration push/pull interaction: This is the fully saturated interaction between a push factor out of an origin office into a destination (the leave-out total number of staffers who ever move between origin and destination) and a pull factor into the destination (the number of staffers who leave the office at time t).

Both of these instruments are arguably uncorrelated with the moving staffer's productivity, potentially satisfying the exclusion restriction. One may be concerned that the turnover instrument may directly influence legislative productivity. We argue that this is unlikely to be substantial for multiple reasons: first, the office fixed effect captures any persistent differences between Representatives who win versus lose re-election. Second, the Representative that loses an election at $t-1$ is no longer in the data at time t, so a violation of the exclusion restriction would have to operate through an anticipated loss affecting legislative productivity prior to the election. However, the incumbent advantage in the US Congress is large, mitigating this potential channel.

We begin by running a standard (linear) first stage, at the staffer-destination office-quarter level. Results are displayed in Table A.3. We note that the most predictive covariate is, unsurprisingly, the destination and origin office being the same. The instruments are all positive and statistically significant, working in the direction that we would expect. Results are robust to the inclusion of various fixed effects, and the $F-$stat is extremely large. Unfortunately, running 2SLS produces extremely wide confidence intervals for staffer shares (the 95% CI is roughly 300% wide). This is because we are simultaneously instrumenting for every single staffer's moves, meaning given I staffers, we have $4I$ instruments for I endogenous variables.

We address this issue by adopting and constructing a stronger first stage with machine learning methods. Instead of predicting the first stage using OLS, we instead use an XGBoosted random forest.[3] In Table A.4, we show that this meaningfully improves the predictive power of the instruments. Column 1 shows the root mean square error and accuracy of the linear first stage, both in and out of the sample. Although the OLS accuracy appears high, recall that a staffer can only be in one destination office at any period; thus, a model that always guesses 0 as the outcome would achieve an accuracy of $\frac{1}{436} = 99.77\%$.[4] Thus, the random forest model in Column 2 performs meaningfully better than the linear model. The upside of the random forest is that it has better predictive power and a stronger first stage. The downside is that it is less interpretable: it is difficult to be sure that the instruments are being used by the model, rather than other features provided to it. We provide suggestive evidence that this is not the case in Column 3, where we show that re-estimating the random forest without instruments provides a worse fit both in and out of sample. This suggests that the model is effectively using information provided by the instruments. To formally test for this, we conduct a permutation test. We randomly permute the instruments across observations and re-estimate the random forest model 1000 times. In every single case, out-of-sample performance using the true instruments beats the model using permuted instruments, establishing that the predictive power of

[3] Dikkala et al. (2020) provides a theoretical justification for non-parametric IV and, in particular, finds that the random forest performs best when compared against a number of other non-parametric estimators.

[4] There are 435 Representatives in the House at any given point in time, plus the outside option of not being in Congress.

the random forest first stage is at least partially driven by the instruments.

To complete the IV estimation, we plug in first-stage estimates from the random forest into the design matrix and re-estimate the linear decomposition following the process described above. The one technical note is that the Bayesian bootstrap must additionally account for the first stage. To keep inference conservative, we independently draw bootstrap weights across the first and second stages and take their product.

Appendix B.2 Bayesian estimation

Hamiltonian Monte Carlo We make use of Bayesian methods for inference. Specifically, we utilize Hamiltonian Monte Carlo (HMC) (Gelman et al., 1995), a Markov Chain Monte Carlo (MCMC) method. Traditional Gibbs and Metropolis algorithms exhibit random walk behavior, causing slow progress as they zigzag through the target distribution. While re-parameterization and refined jumping rules can help, this inefficiency persists in complex, high-dimensional models. HMC enhances the Metropolis algorithm by introducing “momentum" variables, enabling iterations to cover more distance in the parameter space and mix more efficiently. Specifically, for each parameter θ_j in the target space, a corresponding “momentum" variable ρ_j is added. Both θ and ρ are updated together using a modified Metropolis algorithm, where the jumping distribution for θ is influenced by ρ. Each HMC iteration involves multiple steps, during which the position and momentum evolve according to rules that mimic physical motion, allowing for the algorithm to traverse the parameter space more quickly, including turning corners to maintain the trajectory's total “energy".

Specifically, HMC augments the posterior density $p(\theta|y)$ with an independent distribution $p(\rho)$ for “momentum" variables, creating a joint distribution $p(\theta, \rho|y) = p(\rho)p(\theta|y)$. We simulate this joint distribution but focus only on the θ values. Again, the momentum variable ρ helps the algorithm move more efficiently through the parameter space. HMC requires the gradient of the log-posterior density, which must be calculated analytically for efficiency. The momentum ρ is usually assigned a multivariate normal distribution with mean 0 and a covariance matrix M that can be conditioned via warm-up. This setup allows the components of ρ to be independent.

An HMC iteration involves three main steps:

1. Updating Momentum: Draw ρ from its prior distribution $N(0, M)$.
2. Leapfrog Steps: Update θ and ρ through a series of “leapfrog" steps:
    - Half-step for ρ: Update ρ using the gradient of the log-posterior.
    $$\rho \leftarrow \rho + \frac{1}{2}\epsilon \frac{d \log p(\theta|y)}{d\theta}$$
    - Full-step for θ: Update θ using the momentum ρ.
    $$\theta \leftarrow \theta + \epsilon M^{-1}\rho$$
    - Another Half-step for ρ: Update ρ again using the gradient.
    $$\rho \leftarrow \rho + \frac{1}{2}\epsilon \frac{d \log p(\theta|y)}{d\theta}$$

3. Accept/Reject Step: Compute the acceptance ratio r based on the joint densities before and after the leapfrog steps. Accept the new θ with probability $\min(r, 1)$; otherwise, retain the previous θ.

This process is repeated until the MCMC algorithm converges, typically assessed by diagnostics like the R-hat statistic and effective sample size.

For challenging HMC problems, it's ideal for tuning parameters to adapt as the algorithm explores the posterior distribution. This involves scaling the mass matrix M to the local curvature of the log density, adjusting the step size ϵ to be smaller in high-curvature areas, and setting the number of steps to be sufficient for effective exploration without excessive circling. Extensions of HMC, like the no-U-turn sampler (NUTS), achieve this adaptive tuning while maintaining detailed balance. In NUTS, the number of steps is determined adaptively during each iteration, continuing until the trajectory reverses direction. This ensures the trajectory explores the parameter space efficiently. NUTS also includes a procedure to adaptively set M and ϵ during a warm-up phase, which is then fixed for the remaining iterations. We use Stan to run HMC with NUTS. Although tuning HMC can be tricky, NUTS simplifies this by eliminating the need to predefine the number of steps.

Our baseline model uses four independent chains, each run for 2000 iterations (with the first 1000 serving as warm-up/burn-in). We typically set δ (the target average proposal acceptance probability) to 0.95,[5] and the seed for sampling to 1. Our baseline prior for both μ_γ, μ_α is N(0,10), and our baseline prior for $\sigma_\gamma, \sigma_\alpha, \sigma_\tau, \sigma$ is N(0,25).

Appendix B.3 Comparison with other approaches and simulation

In this section, we provide details to the discussion of other leading approaches in the economics literature contained in Section 4.

Chan (2021) uses an empirical Bayes procedure to estimate doctor productivity in a setting where junior and senior physicians are paired together in teams of two. He uses a discrete type space, allowing for 6 types of juniors and 9 types of seniors (based on tenure), and estimates a separate variance for each pair of types. Extending this approach to our setting would require estimating a separate variance for each feasible combination of types in a team (6 types of staffers among the 18 possible "juniors" in a team, and 9 types for the senior, the member of Congress). Assuming symmetry across individuals in a team requires the estimation of $9 * \binom{24}{6} = 1.2$ million parameters; without symmetry, this is $6^{18} * 9 = 1$ quadrillion parameters. Either assumption is infeasible in our context where we lack this many observations—many parameters will remain unidentified.

Bonhomme (2021) estimates researcher productivity using both a linear and non-linear model. His linear model is similar to ours and estimated similarly; our contribution is to focus on a mean decomposition, which also allows us to estimate two novel dynamic specifications. For his nonlinear model, he uses a variational Bayes (ELBO) procedure and restricts attention to teams no larger than two. He estimates his model with K types of workers (using $K = 4$ at baseline). Evaluating the ELBO objective function requires

[5] This is higher than the default used in packages such as RStan, as it produces results that are more robust when posteriors have high curvature, at the cost of computational speed.

a computing a nested sum containing (1) the number of individuals, (2) the number of teams, and (3) the number of types for each individual on the team. For the simplest case of a single office containing 18 staffers, this is $18 * 4^{18} = 1$ trillion operations. This is a lower bound as there are many unique teams in the data (on the order of 10,000), although not all of them hire the maximal 18 staffers. Actual estimation typically requires many calls to the objective function, and is computationally infeasible.

Ahmadpoor and Jones (2019) pose a nonlinear least squares problem that they solve using simple gradient descent. In particular, they pick an informed initialization, run gradient descent for one hundred thousand iterations, and conclude the procedure. To compare this method against our Bayesian procedure, we first generate 50 simulated datasets that use the same history of moves as our primary sample, and simulated staffer and Representative effects using the same distributions as the baseline estimates presented in Table 3, Panel B. We then compare estimates from both procedures. In Figure 8, we plot estimated values of staffer substitutability (ρ_{staff}) and mean office productivity ($\exp(\mu_\gamma + \frac{\sigma_\gamma}{2})$), alongside the true parameter value. While the Bayesian method is randomly initialized, the nonlinear least squares method is initialized with the oracle solution. Despite this asymmetry, we find that the Bayesian method does better at estimating these parameters.

Bergeron et al. (2022) adopts a linear model similar to ours. However, because output is non-parametrically estimated as a combination of discrete types, this would still result in over 200,000 parameters to estimate in our context, roughly an order of magnitude above the number of observations that we have.

In addition to these papers, we are also aware of Freund (2024), Iranzo et al. (2008), and Weidmann et al. (2024), who all study the contribution of individuals in team environments. Each of these papers directly estimates an individual's effect, and uses these estimates to study their impact on teams. Iranzo et al. (2008) estimate individual effects through a standard AKM wage regression, and plug these effects into their firm production estimation procedure. Freund (2024) estimates individual effects through an augmented AKM regression that takes the team into account, but the outcome data are ultimately still individual-level wages. Weidmann et al. (2024) directly measures individual effects in the lab before randomly pairing individuals into teams.

Finally, we discuss the simulation procedure that we use to validate our estimates. For the dynamic models, we simulate staffer and office effects using the same distributions as the baseline estimates presented in Table 3, Panel A. We then construct a panel of movers that matches the main mover sample along the following moments: the average and standard deviation for number of quarters spent in an office, the average and standard deviation for number of moves between offices, and the average number of quarters that a Representative spends in Congress. This leads to an office share of differences between above and below median offices of 73%, and an office share of differences between destination and origin offices of 8 to 12% (depending on whether odd or even quarters are used). In estimation, we find an office share of differences between above and below median offices of 72%, with a 95% CI of [0.713, 0.733]. We plot the dynamic estimate of the office share of differences between destination and origin offices in Figure A.14, Panels A and B, and find them consistent with the true values. We also plot predicted staffer, office, and other team-member moves in Figure A.14, Panels C and D, and find results

consistent with the discussion of Equation 7. For our Bayesian procedure, we note that, by design, inference on the posterior is correct. Nevertheless, we evaluate the nested CES estimation procedure and find that the true parameter values fall within the 95% credible interval.

Appendix B.4 Estimating bill proposal locations

The reported ideological position of a bill in DW-Nominate represents the estimated cut-point that separates a bill's proposed ideological position from the status quo. Although these cut-points may be used as proxies for a bill's proposed position, the cut-points are generally only unbiased under assumptions such as a symmetrically distributed status quo (Vander Wielen and Vander Wielen, 2020). To estimate a bill's proposal location, we follow the approach in Woon (2008), which uses co-sponsorship data to identify the location of a proposal.

Specifically, Woon (2008) demonstrates that under a threshold model of co-sponsorship (where Representatives only co-sponsor bills sufficiently close to them ideologically), the proposal location of a bill can be estimated by first fitting the following probit model for each bill:

$$Pr(Cosponsor_j) = \Phi(\xi_2 \iota_j^2 + \xi_1 \iota_j + \xi_0)$$

where ι_j is the DW-Nominate measure of a Representative j's ideology and $Cosponsor_j$ is an indicator for whether Representative j cosponsored the bill. The implied proposal location is simply the maximum of the estimated quadratic function: $\hat{\iota}_b = -\frac{\hat{\xi}_1}{2\hat{\xi}_2}$.

Since we are not interested in individual bills, but in the average ideology of all bills produced by an office j in each quarter t, we pool estimation for all bills at the office-quarter level to estimate an office's ideal proposal location ι_{jt}.[6] This serves as an alternate outcome that may be used to estimate Equation 11.

Appendix C Additional results

Appendix C.1 The returns to experience

We examine the role that experience plays in determining an individual's legislative productivity. We estimate an extension of the model in Equation 1 that allows for experience to flexibly scale up (or down) the productivity of staffers and Representatives based on their experience in Congress. Specifically, we estimate the model:

$$y_{jt} = \sum_i [\alpha_i m_{i,j,t} \tau_{total}(i,t)^{\phi_{staff}}] + \gamma_j (\tau_{total}(j,t)^{\phi_{rep}}) + \tau_t + \varepsilon_{jt} \tag{12}$$

where τ_{total} reflects the total amount of time (in quarters) that a staffer or Representative has spent in Congress, and the exponent ϕ is the parameter of interest.

We present trace plots of the estimated ϕ_{staff} and ϕ_{rep} in Figure A.16. We find that the posterior mean for ϕ_{rep} is 0.17, implying that Representatives with 1 full session of experi-

[6] Following Woon (2008), we drop all office-quarters when $\xi_2 > 0$ (to ensure we have found a maximum) or the standard error of the estimated ι_{jt} is greater than 2.

ence in Congress (8 quarters) are 42% more effective than a freshman with no experience, and that Representatives with 10 full sessions of experience in Congress (80 quarters) are 210% more effective than a freshman with no experience. On the other hand, we are unable to reject the null of staffer experience being irrelevant, $\phi_{\text{staff}} = 0$. We note that although our estimates are imprecise (a 95% confidence interval would just barely exclude gains comparable to Representatives, $\phi_{\text{staff}} > 0.17$), the lack of a detectable effect is consistent with Crosson et al. (2020), who find that total legislative staff experience does not correlate with lawmaking effectiveness.[7]

Staffer skill and exit The revolving door between staffing and lobbying has received substantial attention (Bertrand et al., 2014). Given their low pay, a concern is that skilled staffers may disproportionately exit Congress for other jobs.[8] To test this, we regress posterior means of $\log(\alpha_i)$ against the length of each staffer's career in Congress (in quarters). Using our baseline estimates of α_i, we find a coefficient (SE) of -0.412 (0.139). Using estimates of α_i from our experience model, we find a coefficient (SE) of -0.485 (0.143). These suggest that productive staffers spend less time in Congress. A staffer who is twice as productive as another is expected to spend 1.5 fewer months in Congress.

[7]Specifically, Crosson et al. (2020) finds that the average effect of legislative staff experience is indistinguishable from 0, but that experienced staff positively impact legislative productivity for committee chairs and inexperienced members of Congress.

[8]See, for instance, `https://sunlightfoundation.com/2010/12/21/keeping-congress-competent-staff-pay-turnover-and-what-it-means-for-democracy/`.

Appendix D Tables and Figures

Table A.1: Correlating staffer FEs with each other

FE source:	Bills cospons.	Exit comm.	Pass House	Laws passed
	(1)	(2)	(3)	(4)
OLS: # orig. bills cospons.	1.000			
	(.)			
OLS: # bills becoming law	0.358	1.000		
	(0.039)	(0.000)		
OLS: # bills exit committee	0.528	0.570	1.000	
	(0.037)	(0.040)	(0.000)	
OLS: # bills pass House vote	0.537	0.592	0.985	1.000
	(0.038)	(0.040)	(0.007)	(0.000)

Notes: Regressions are at the staffer level. FEs are sourced from an OLS estimate of Equation 1. Standard errors are robust

Table A.2: Misspecified AKM designs, staffer share of differences

Outcome:	Bills cospons.	Laws passed	Exit comm.	Pass House	Important bills
	(1)	(2)	(3)	(4)	(5)
Panel A: baseline, team-based AKM model					
Staffer share	0.489	0.240	0.212	0.395	0.141
	[0.416,0.591]	[0.055,0.458]	[0.028,0.416]	[0.152,0.575]	[-0.143,0.357]
Panel B: misspecified, standard AKM model (all staffers)					
Staffer share	0.233	0.034	0.022	-0.019	0.066
	[0.157,0.302]	[-0.083,0.130]	[-0.090,0.150]	[-0.143,0.116]	[-0.051,0.176]
Panel C: misspecified, standard AKM model (only Chiefs of Staff)					
Staffer share	0.215	0.055	0.080	0.047	0.072
	[0.188,0.248]	[0.020,0.145]	[0.010,0.167]	[-0.036,0.104]	[-0.014,0.154]
Panel D: misspecified, standard AKM model (only Legislative Directors)					
Staffer share	0.218	0.049	0.050	0.047	0.006
	[0.171,0.165]	[-0.067,0.099]	[-0.005,0.109]	[-0.017,0.197]	[-0.140,0.161]

Notes: This table presents the staffer share of differences between above and below-median offices across a number of legislative outcomes. In Panel A, we present the baseline estimates from the team-based AKM model as in Figure 1. Panels B-D present misspecified models that run the standard AKM design at the individual level. Panel B includes all movers, meaning an observation is repeated as many times as there are staffers in the office. Panels C-D restrict there to be a single staffer per office. In Panel C, this is the Chief of Staff. In Panel D, this is the Legislative Director. Standard errors are computed via 100 iterations of the Bayesian bootstrap. 95% confidence intervals are displayed beneath point estimates in brackets.

Table A.3: IV — linear first stage

	(1)	(2)	(3)	(4)
Stay in same office	0.807	0.807	0.798	0.798
	(0.002)	(0.002)	(0.002)	(0.002)
Turnover	0.021	0.021	0.021	0.021
	(0.005)	(0.005)	(0.005)	(0.005)
Total moves from orig. to dest.	0.019	0.018	0.019	0.019
	(0.001)	(0.001)	(0.001)	(0.001)
# leaving dest.	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Moves from O-D X # leave dest.	0.006	0.006	0.006	0.006
	(0.000)	(0.000)	(0.000)	(0.000)
Dist. in party share	0.000	0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
N	21018181	21018180	21018180	21018180
F	42942.09	43095.46	36766.30	37176.89
RMSE	0.031	0.031	0.031	0.031
FE	None	Time	+ Orig./Dest.	+ Staffer

Notes: This table presents regression coefficients and standard errors for the first stage of the IV at the staffer-origin office-destination office-quarter level. Column 2 adds time-fixed effects, column 3 adds origin and destination-fixed effects, and column 4 adds staffer-fixed effects. Standard errors are clustered at the staffer-quarter level.

Table A.4: IV — comparing first stages

	OLS	RF	RF (no instruments)	(2) < (3) $p-$val.
	(1)	(2)	(3)	(4)
RMSE (in sample)	0.0310	0.0273	0.0274	-
Accuracy (in sample)	99.894%	99.915%	99.914%	-
RMSE (out of sample)	0.0312	0.0266	0.0280	< 0.001
Accuracy (out of sample)	99.892%	99.920%	99.912%	> 0.999

Notes: This table presents in-sample and out-of-sample root mean square error (RMSE) and accuracy for various IV first stages. Column 1 presents results from OLS (Table A.3, Column 4). Column 2 presents results from an XG-Boosted random forest. Column 3 presents results from the random forest without instruments. $p-$values in column 4 are computed using a permutation test (permuting instrument rows then re-fitting the random forest) with 1000 draws.

Table A.5: Bayesian model, HMC NUTS sampler properties

Property	Min.	25%	Median	Mean	75%	Max
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: baseline model						
Accept stat.	0.00	0.84	0.94	0.85	0.98	1.00
Step size	0	0.001	0.001	0.002	0.002	14.
Tree depth	0.0	8.0	8.0	7.9	8.0	10.0
N leapfrog	1	255	255	296	255	1023
Divergent	0.000	0.000	0.000	0.052	0.000	1.000
Panel B: CES model						
Accept stat.	0.00	0.94	0.98	0.94	1.00	1.00
Step size	0	0.0015	0.0079	0.0126	0.0089	10.95
Tree depth	0.0	9.0	9.0	8.9	10.0	10.0
N leapfrog	1	511	511	662	1023	1023
Divergent	0.000	0.000	0.000	0.016	0.000	1.000

Notes: This table presents sampler properties from the baseline model (Panel A) and CES model (Panel B). Columns 2 through 7 display the minimum, 25th percentile, median, mean, 75th percentile, and maximum respectively. Accept stat. is the probability that a proposal is accepted. Step size controls how far the momentum variable updates. Tree depth is the depth of the binary tree used by NUTS, which determines the number of leapfrog steps taken by the sampler to avoid random walk behavior. Divergent transitions are instances when the Hamiltonian value deviates substantially from the initial value, indicating potential random walk behavior.

Table A.6: Correlation between Bayesian ideology estimates and DW-Nominate scores

	1st Dimension		2nd Dimension	
	Dem.	Rep.	Dem.	Rep.
	(1)	(2)	(3)	(4)
DW-Nominate score	-0.0144	0.0184**	0.1468*	0.0296**
	(0.0756)	(0.0089)	(0.0845)	(0.0126)

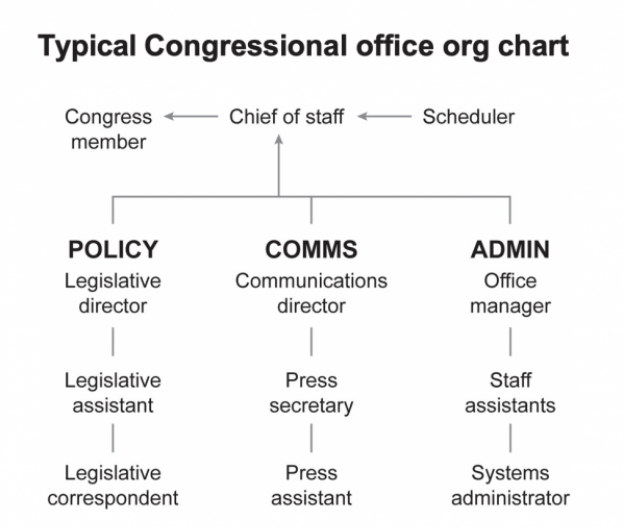
Notes: This table presents regressions at the Representative level, where the outcome is the posterior mean of a Representative's ideology when writing bills, and the independent variable is the Representative's DW-Nominate score from roll call votes. Columns 1 and 2 present results from the 1st dimension of DW-Nominate, and columns 3 and 4 from the 2nd dimension. Odd rows restrict the sample to Democrats, even rows to Republicans. Standard errors are robust. * significant at 10% ** significant at 5% *** significant at 1%.

Table A.7: Bayesian estimates of staffer shares, robustness

Parameter	Mean	2.5%	97.5%	$\tilde{n}$	$\hat{R}$
(1)	(2)	(3)	(4)	(5)	(6)
Panel A: baseline model					
Staffer share	0.37	0.34	0.41	1071	1.00
Panel B.1: baseline model, top vs. bottom tercile offices					
Staffer share	0.37	0.34	0.41	999	1.00
Panel B.2: baseline model, top vs. bottom quartile offices					
Staffer share	0.38	0.34	0.41	953	1.00
Panel C.1: alternate distribution, normal					
Staffer share	0.43	0.38	0.48	3424	1.00
Panel C.2: alternate distribution, Fréchet					
Staffer share	0.39	0.35	0.42	1062	1.00
Panel C.3: alternate distribution, mixture of normals					
Staffer share	0.43	0.39	0.47	4	1.39
Panel D.1: alternate outcomes, bills exiting committee					
Staffer share	0.44	0.38	0.49	1076	1.00
Panel D.2: alternate outcomes, bills passing House vote					
Staffer share	0.44	0.38	0.49	969	1.00
Panel D.3: alternate outcomes, bills becoming law					
Staffer share	0.53	0.47	0.59	1203	1.00
Panel D.4: alternate outcomes, important bills					
Staffer share	0.83	0.76	0.90	504	1.01
Panel E.1: including non-movers					
Staffer share	0.63	0.59	0.69	628	1.01

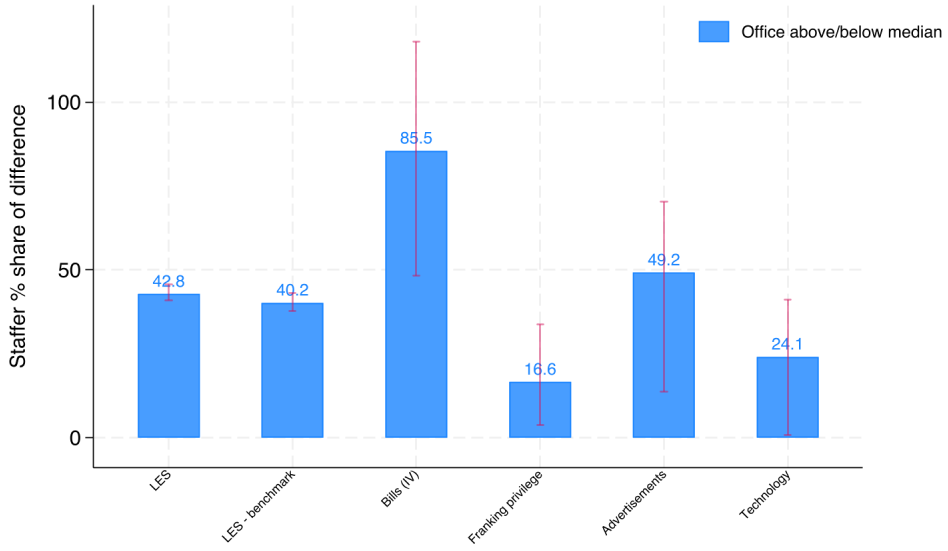
Notes: This table presents posterior means, 2.5th and 97.5th percentiles of staffer shares. Column 5 displays the effective sample size and column 6 displays the $\hat{R}$, a measure of MCMC convergence. Panel A presents the staffer share from the baseline model. Panel B presents estimates from the top vs. bottom tercile/quartile office instead of above vs. below median offices. Panel C presents staffer shares under alternate distributions for both staffers and offices. Panel D presents staffer shares for other legislative outcomes. Panel E presents staffer shares when including non-movers as well.

Figure A.1: Staffer org chart



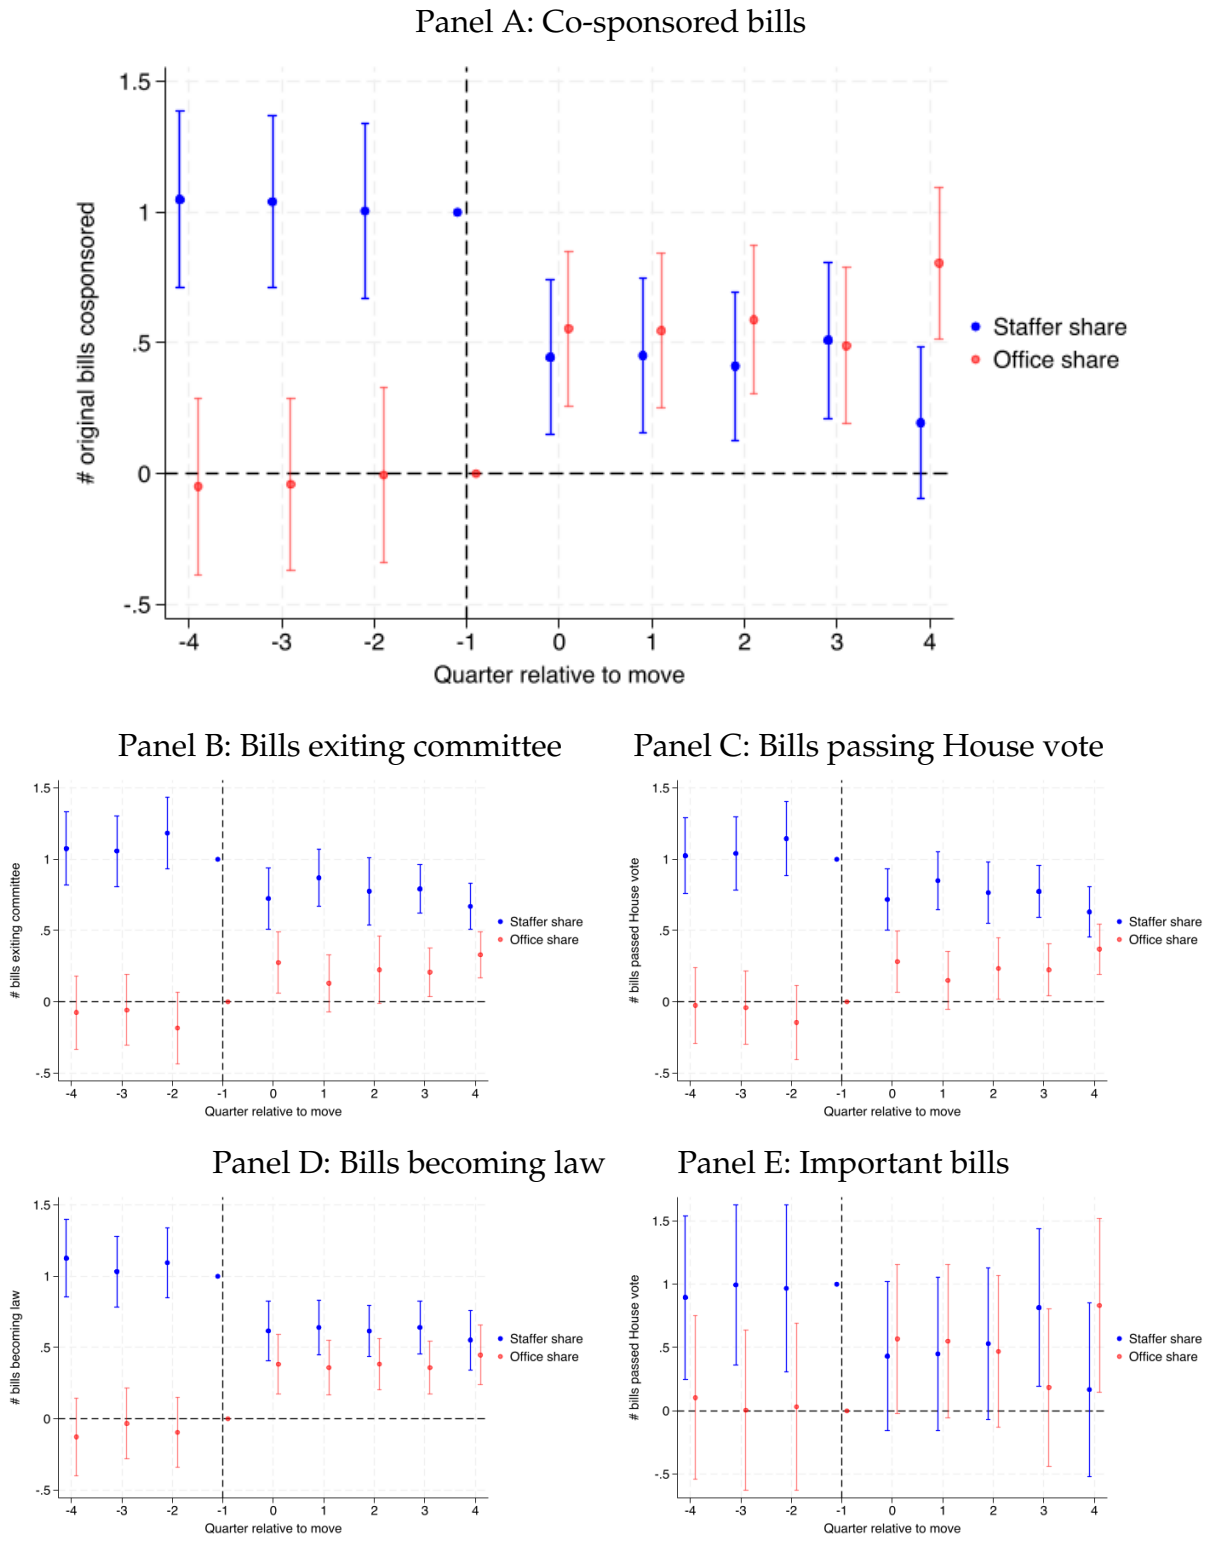
Note: Source `https://www.politicopro.com/blog/congressional-office-org-chart/`

Figure A.2: Staffer share of differences in office spending, linear decomposition



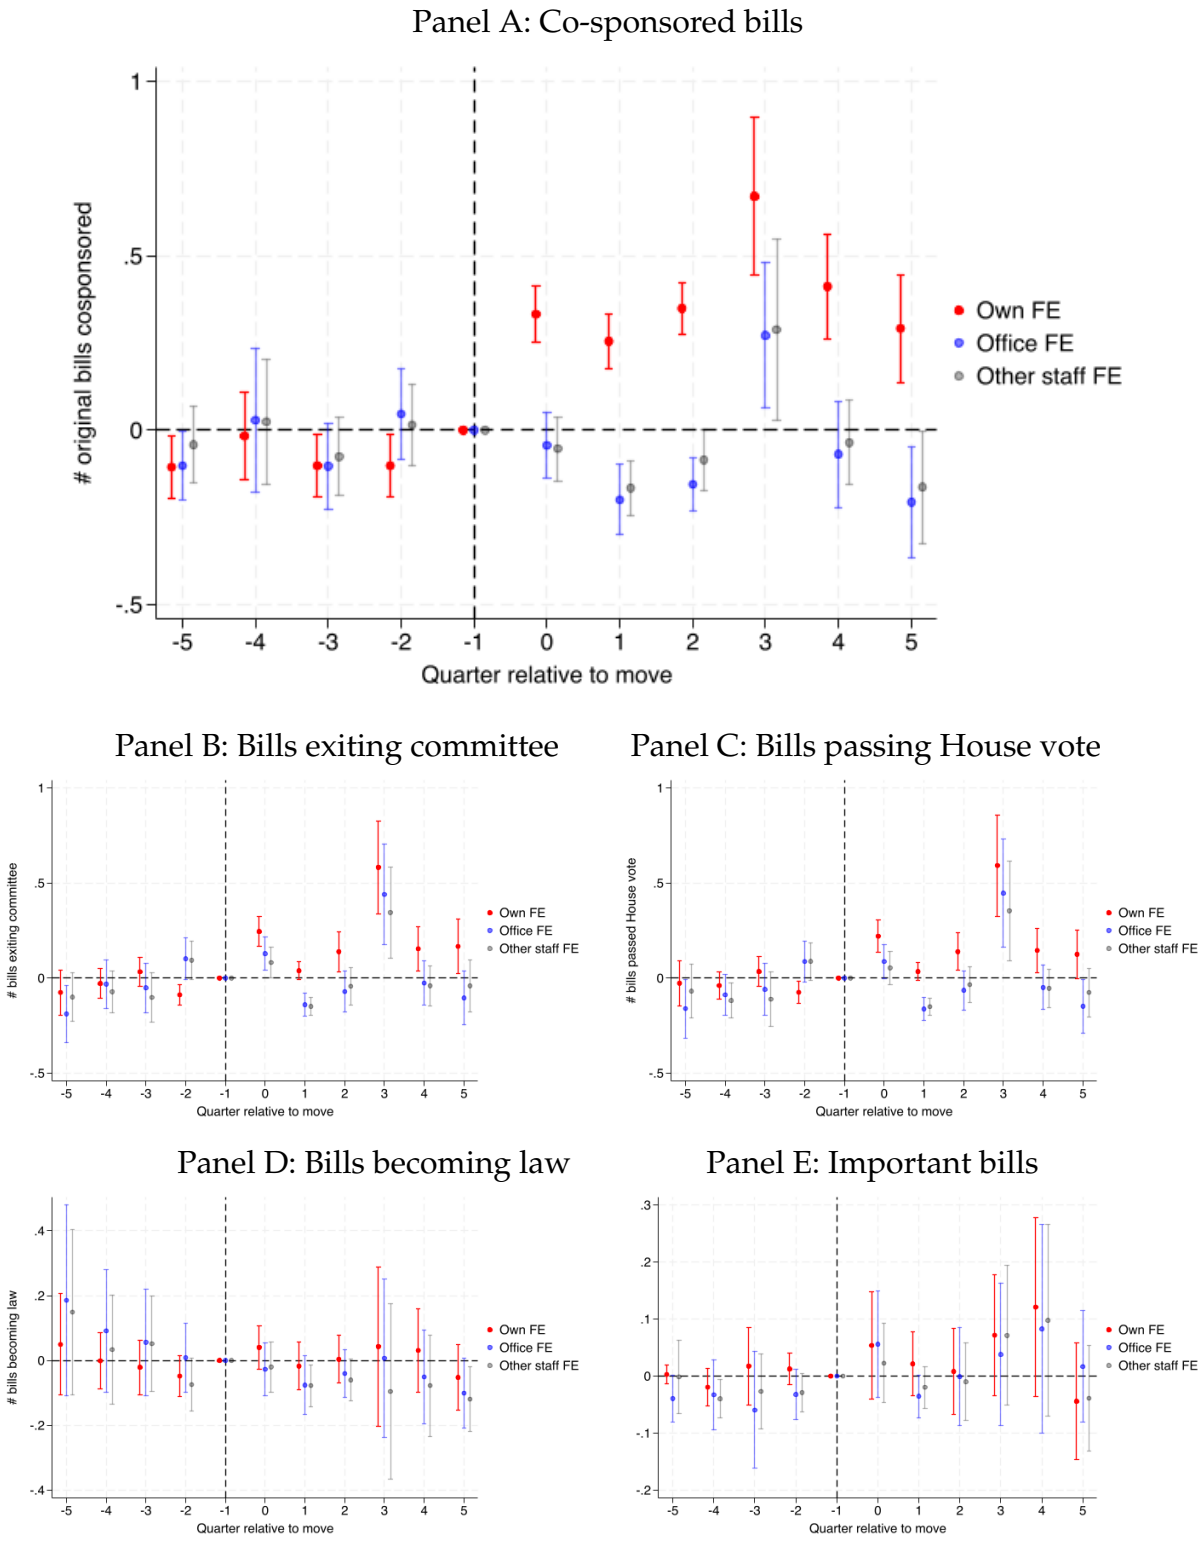
Note: This figure plots the staffer share of differences legislative productivity (first three columns) and office spending (last three columns), as well as 95% confidence intervals. The staffer share is computed from fixed effects estimated in a team-based mover design framework (Equation 1). LES refers to the Legislative Effectiveness Score. The third column utilizes an IV strategy, where the staffer share is computed from instrumented fixed effects estimated in a team-based mover design framework (Equation 1), using a random forest first stage.

Figure A.3: Dynamic estimation of office shares, in-sample (odd quarters)



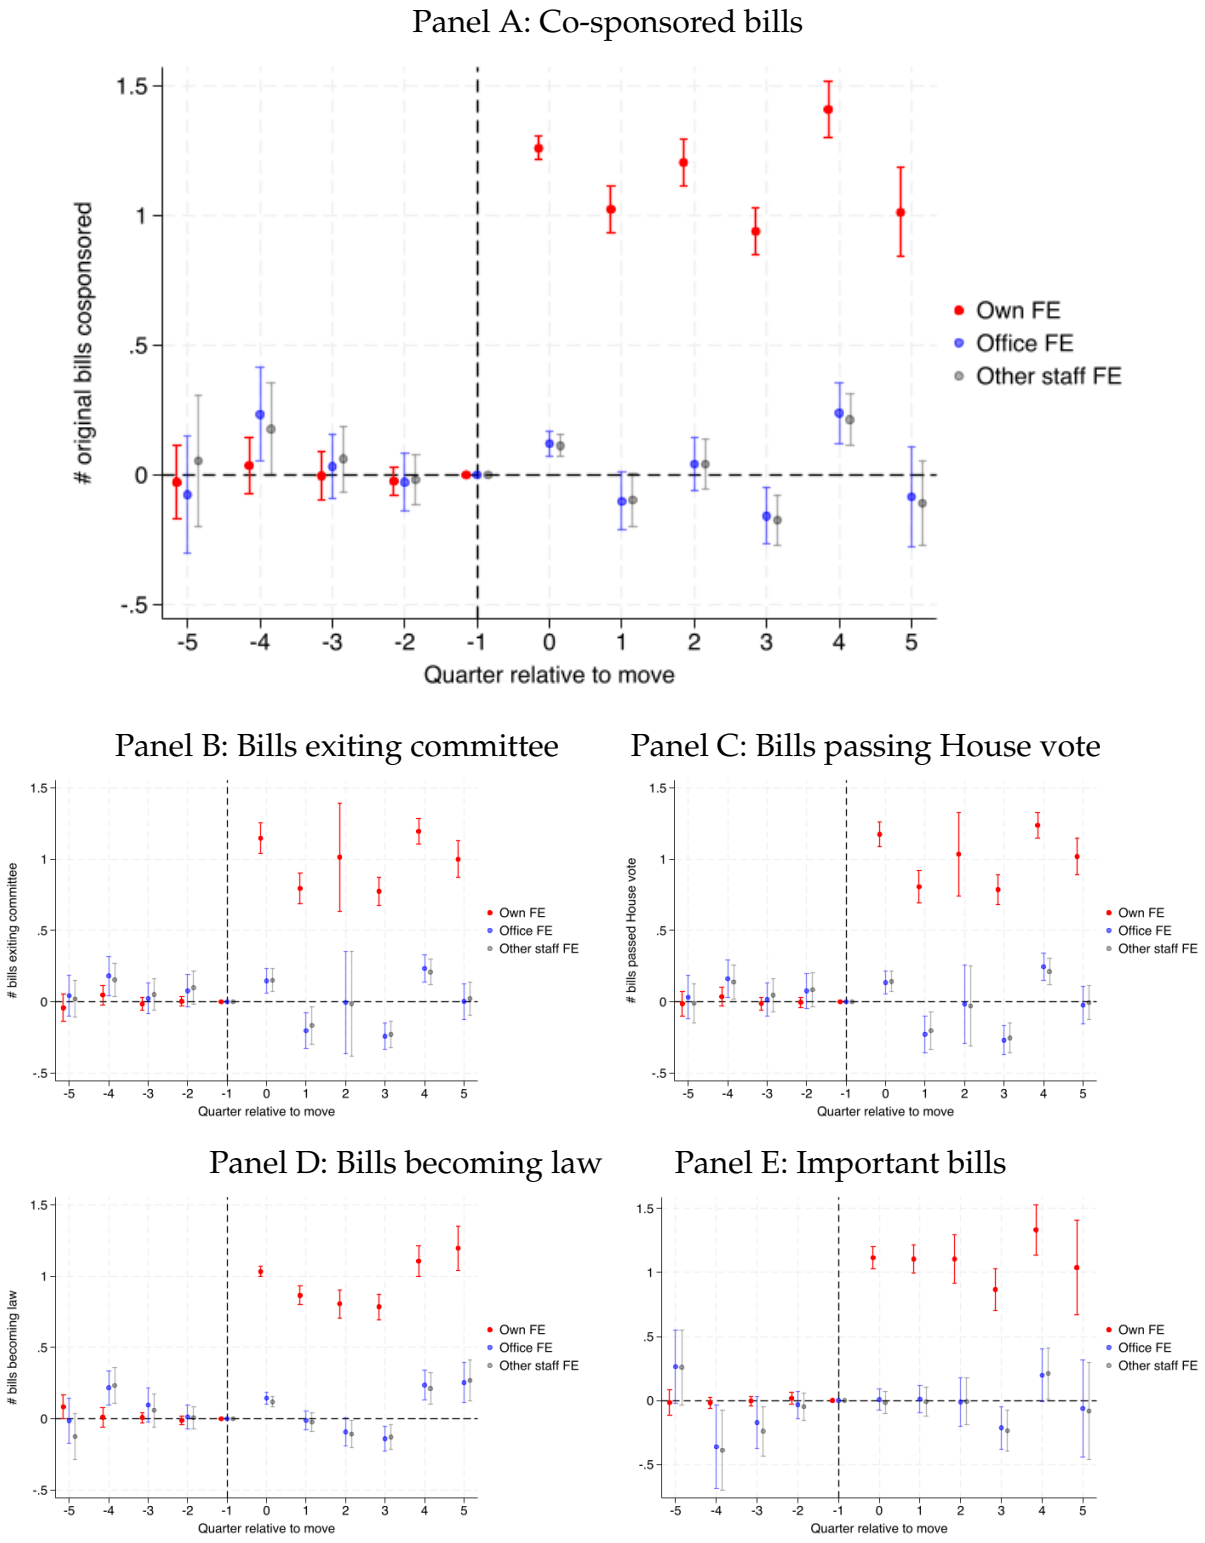
Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 6) at the staffer-quarter level (for odd quarters only). The x-axis shows quarters to a staffer's first move. Displayed coefficients are the office share of differences between the destination and origin office's legislative output. The specific type of legislative output varies by panel. Standard errors are clustered at the office level.

Figure A.4: Dynamic prediction of staffer movement using fixed effects



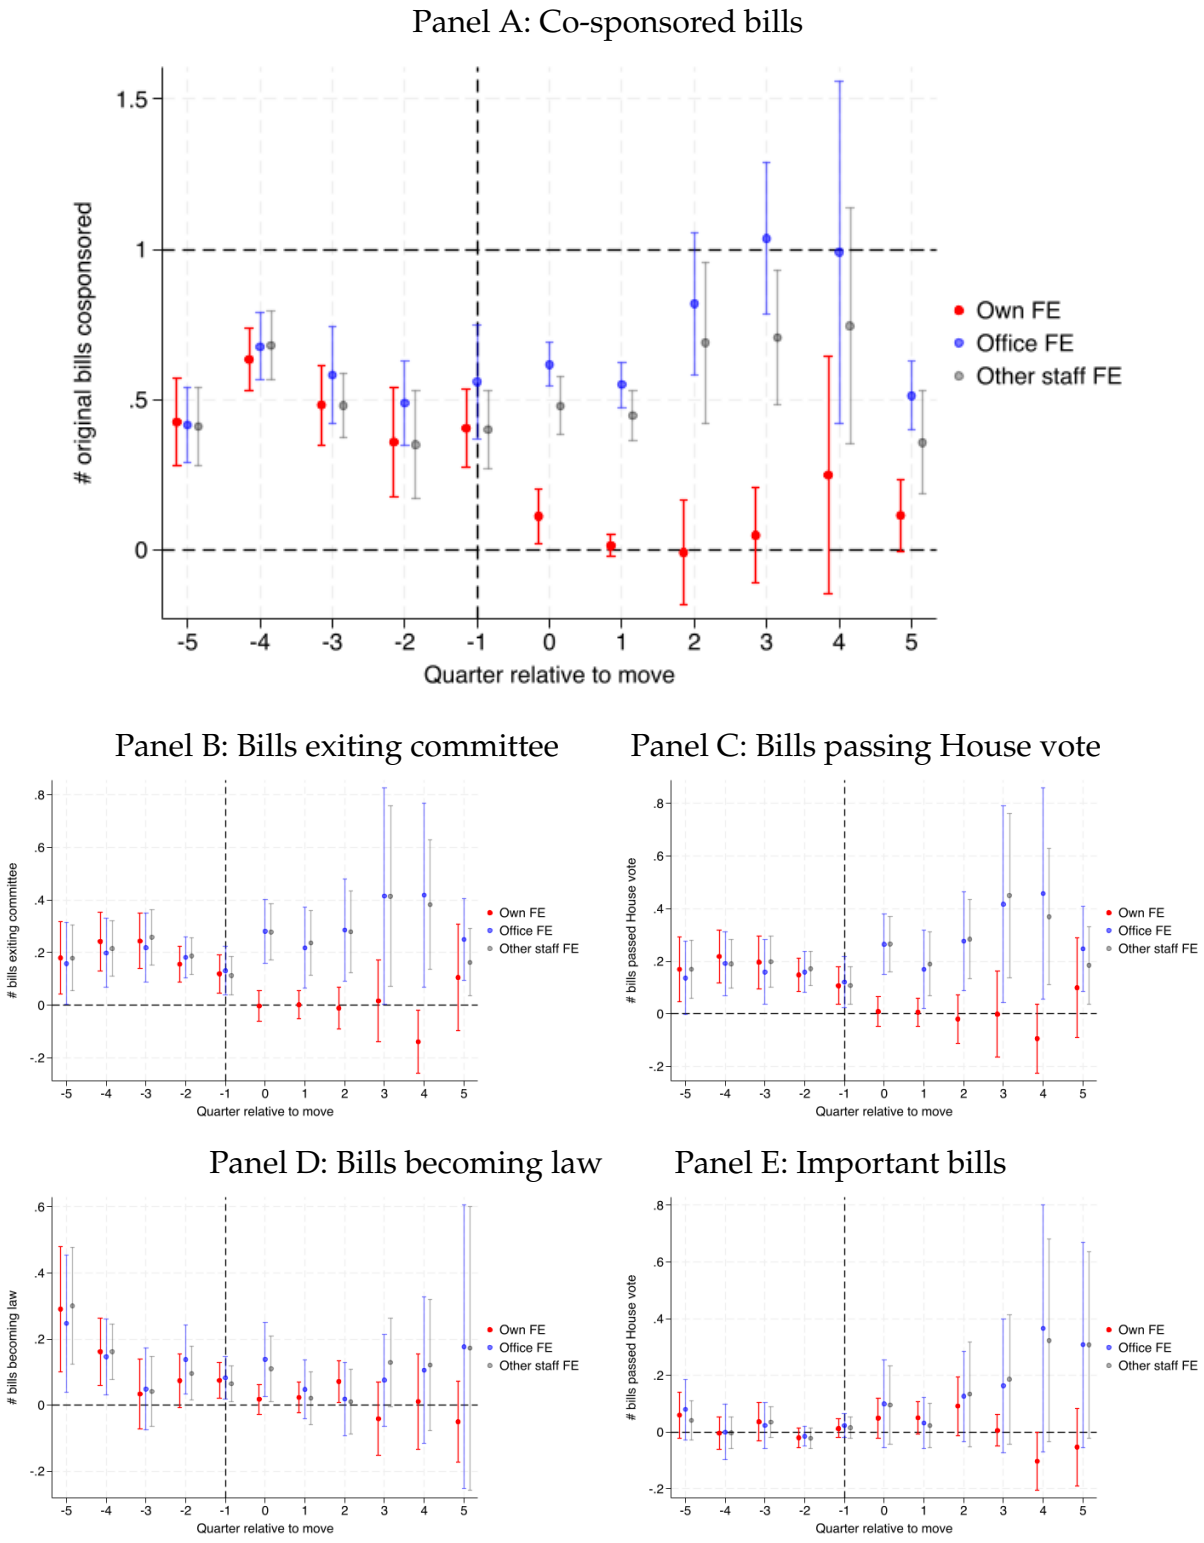
Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 7) at the staffer-quarter level (for even quarters only). The x-axis shows quarters to a staffer's first move. The outcome is a specific type of destination office legislative output that varies by panel. Coefficients on the moving staffer's fixed effect are plotted in red, coefficients on the destination office's fixed effect are plotted in blue, and coefficients on the destination team of staffers' fixed effect are plotted in gray. All fixed effects are estimated out of sample on odd quarters. Standard errors are clustered at the office level.

Figure A.5: Dynamic prediction of staffer movement, in-sample (odd quarters)



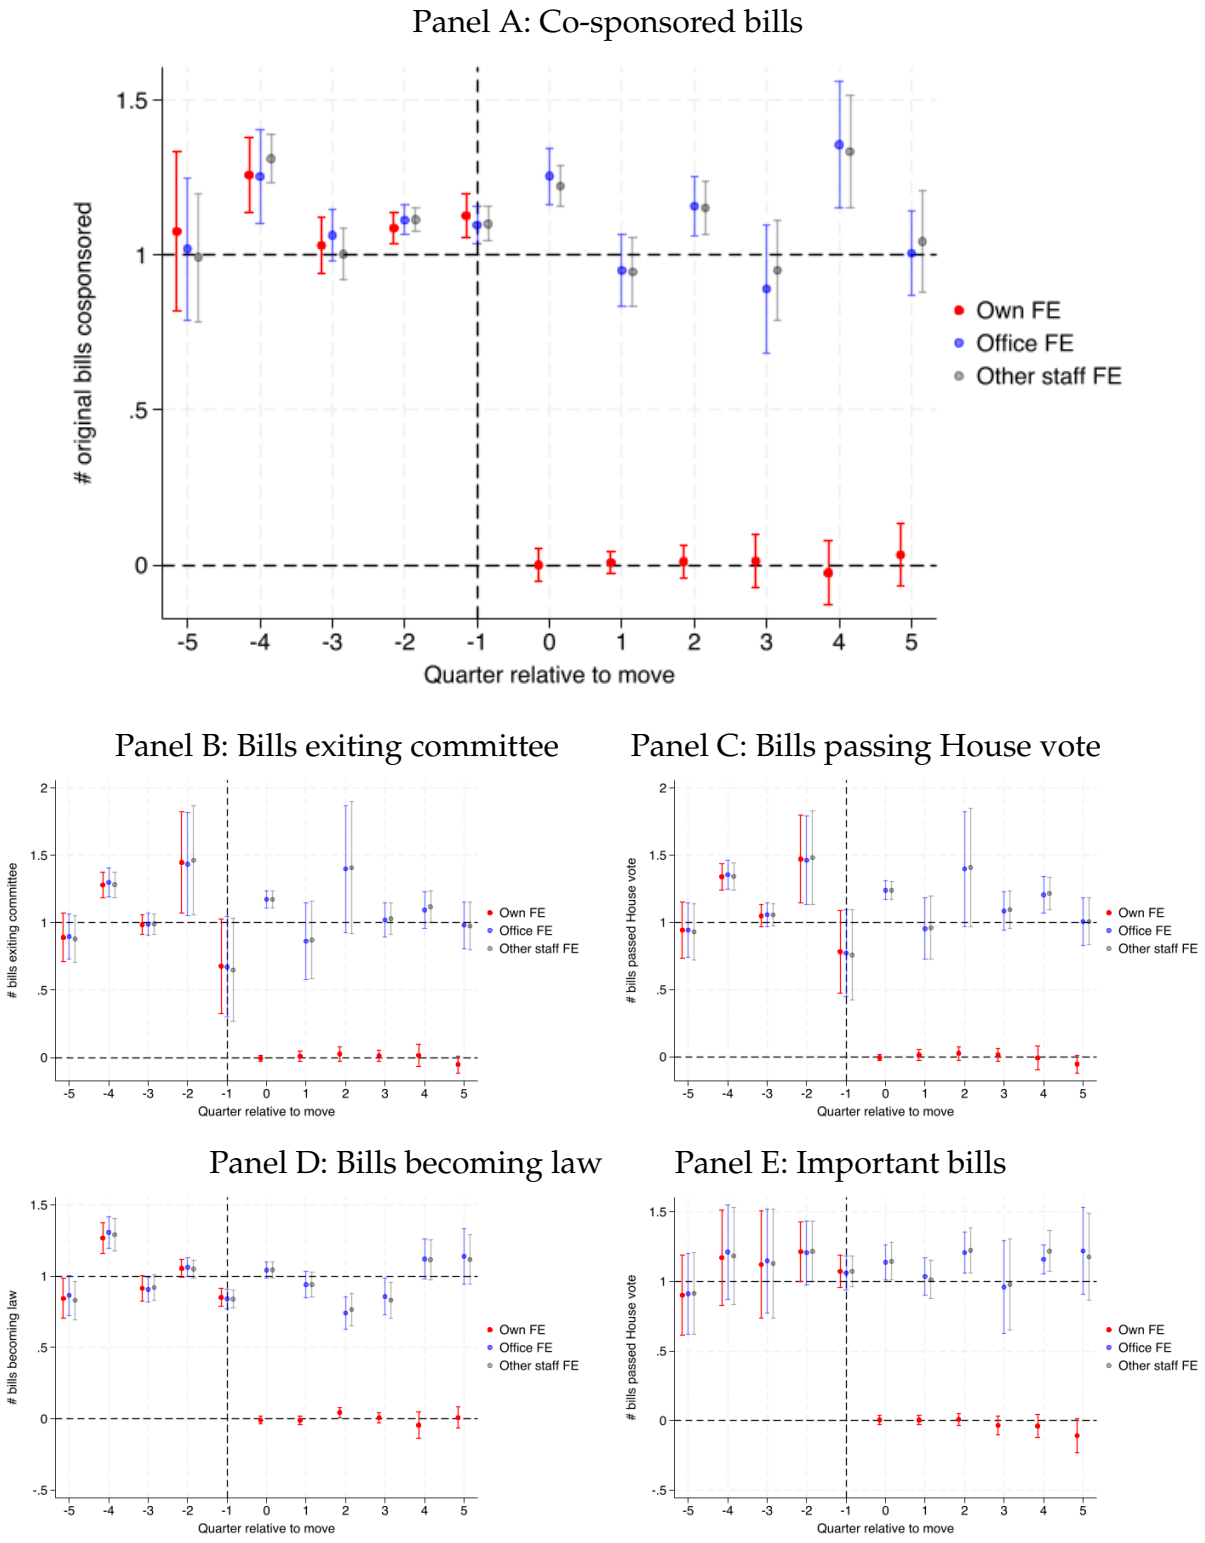
Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 7) at the staffer-quarter level (for odd quarters only). The x-axis shows quarters to a staffer's first move. The outcome is a specific type of destination office legislative output that varies by panel. Coefficients on the moving staffer's fixed effect are plotted in red, coefficients on the destination Representative's fixed effect are plotted in blue, and coefficients on the destination team of staffers' fixed effect are plotted in gray. Standard errors are clustered at the Representative level.

Figure A.6: Dynamic prediction of staffer movement, from origin office



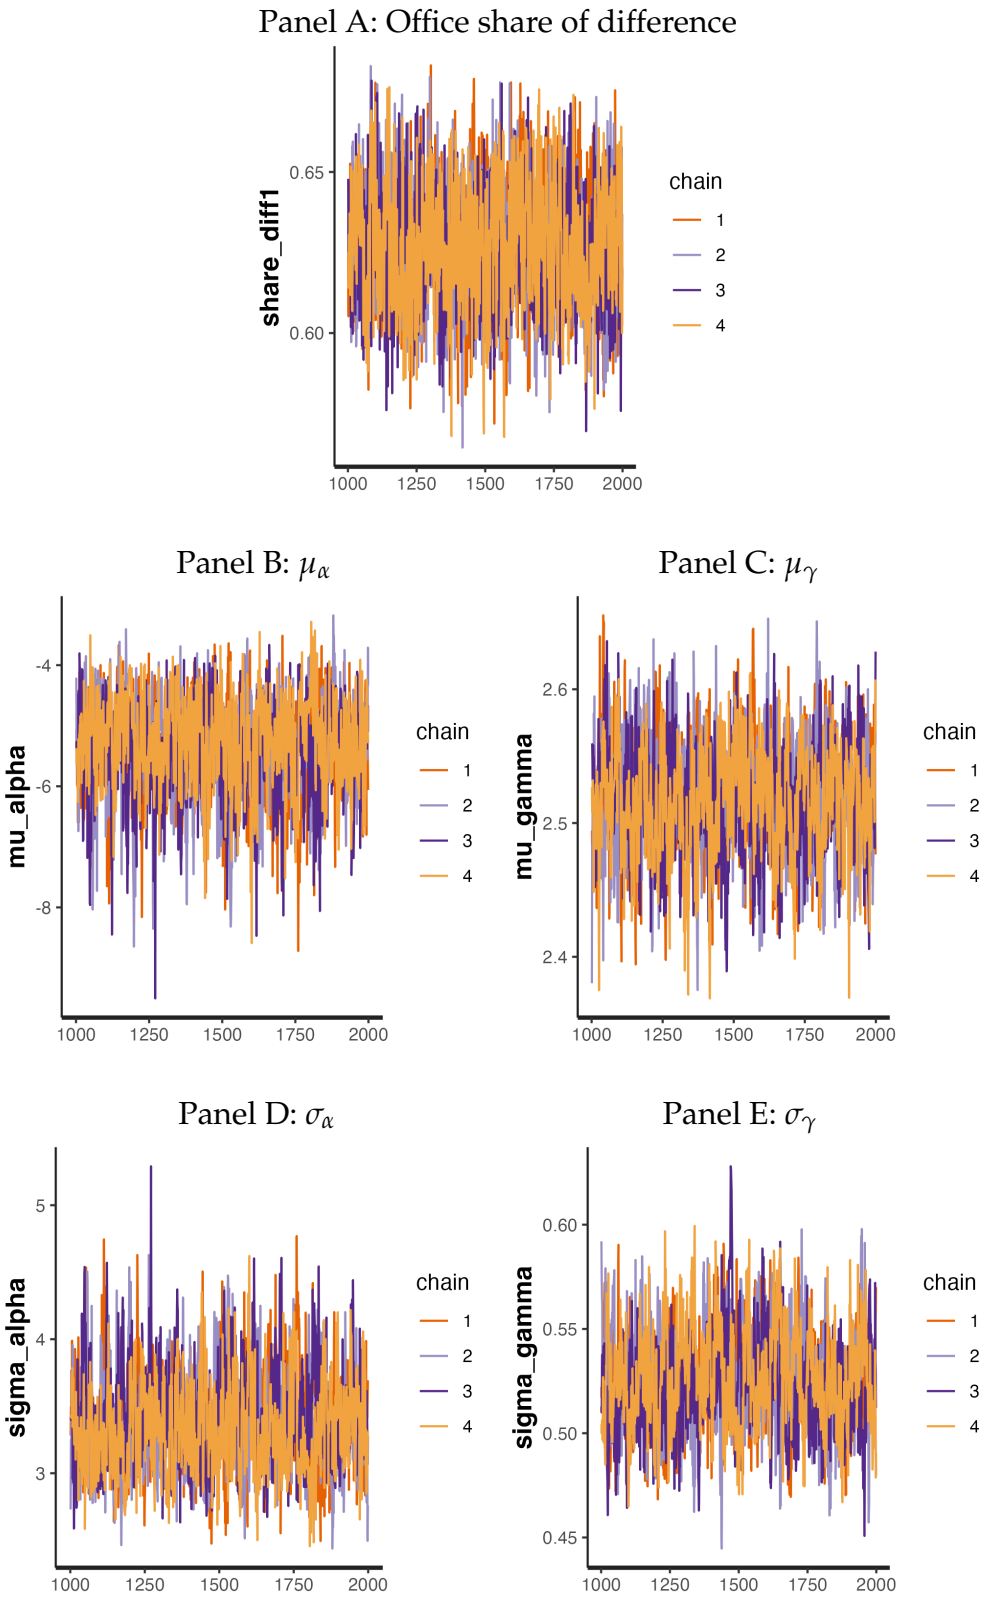
Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 7) at the staffer-quarter level (for even quarters only). The x-axis shows quarters to a staffer's first move. The outcome is a specific type of origin office legislative output that varies by panel. Coefficients on the moving staffer's fixed effect are plotted in red, coefficients on the destination office's fixed effect are plotted in blue, and coefficients on the destination team of staffers' fixed effect are plotted in gray. These fixed effects are estimated out of sample on odd quarters. Standard errors are clustered at the office level.

Figure A.7: Dynamic prediction of staffer movement, from origin office, in-sample (odd quarters)



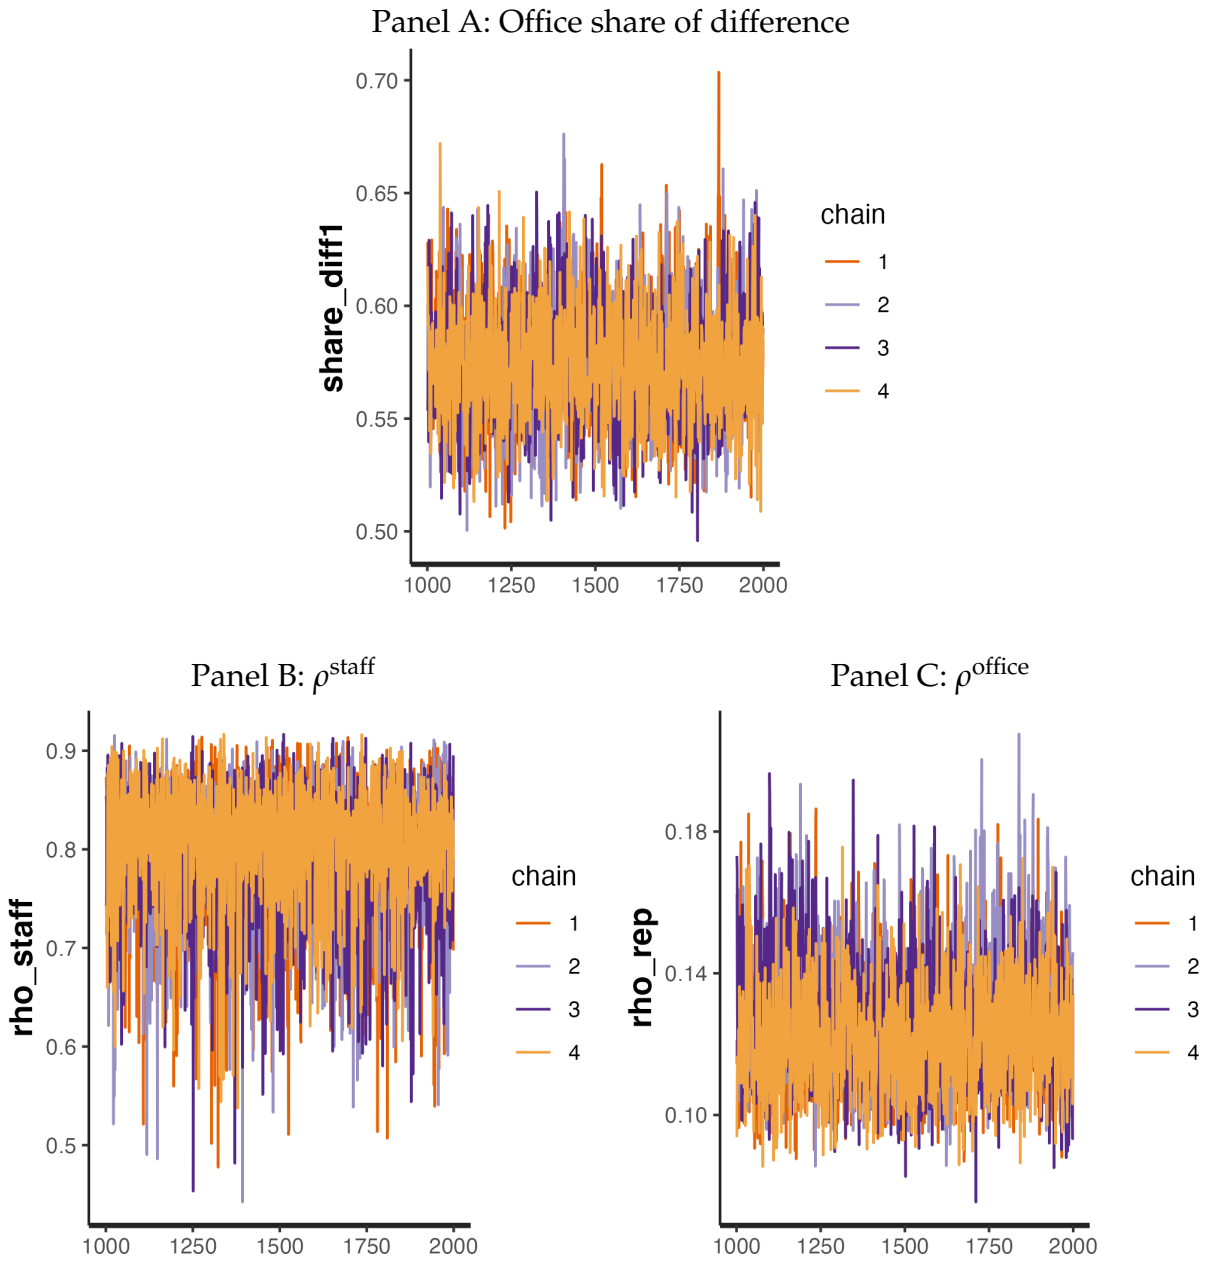
Note: This figure plots coefficients and 95% confidence intervals for an event study specification (Equation 7) at the staffer-quarter level (for odd quarters only). The x-axis shows quarters to a staffer's first move. The outcome is a specific type of origin office legislative output that varies by panel. Coefficients on the moving staffer's fixed effect are plotted in red, coefficients on the destination office's fixed effect are plotted in blue, and coefficients on the destination team of staffers' fixed effect are plotted in gray. Standard errors are clustered at the office level.

Figure A.8: Trace plots for the baseline model, number of bills co-sponsored



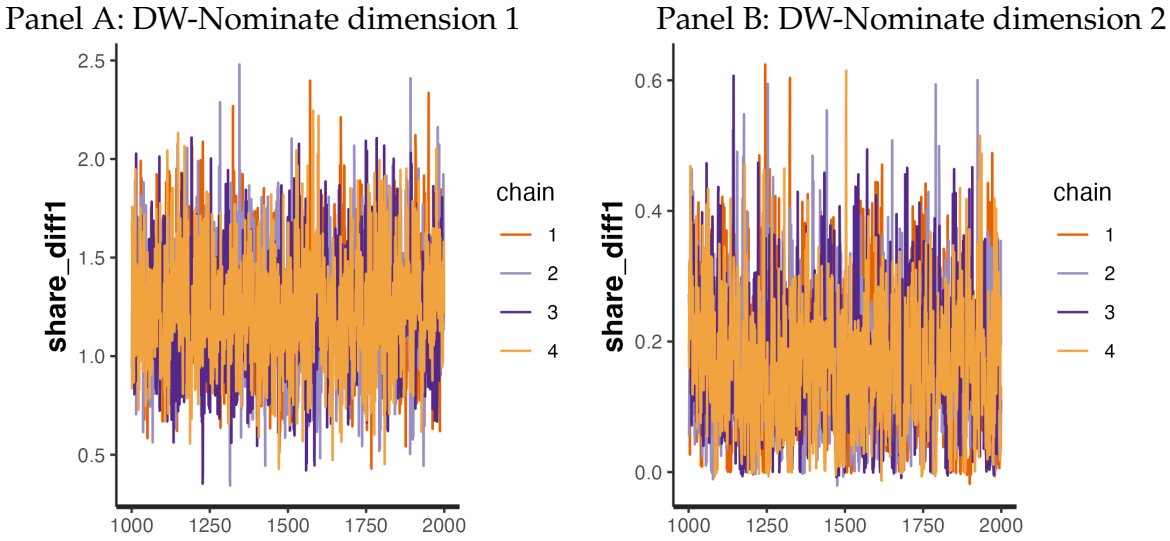
Note: This figure presents trace plots for a Bayesian estimation of the team-based mover design framework (Equation 1). The x-axis represents samples from a posterior drawn via MCMC. Panel A shows the office share of differences between above and below-median offices in bills produced. Panels B and D (C and E) display the staffer (office) log-normal distribution's mean and standard deviation. Values for 4 different chains are displayed.

Figure A.9: Trace plots for CES model



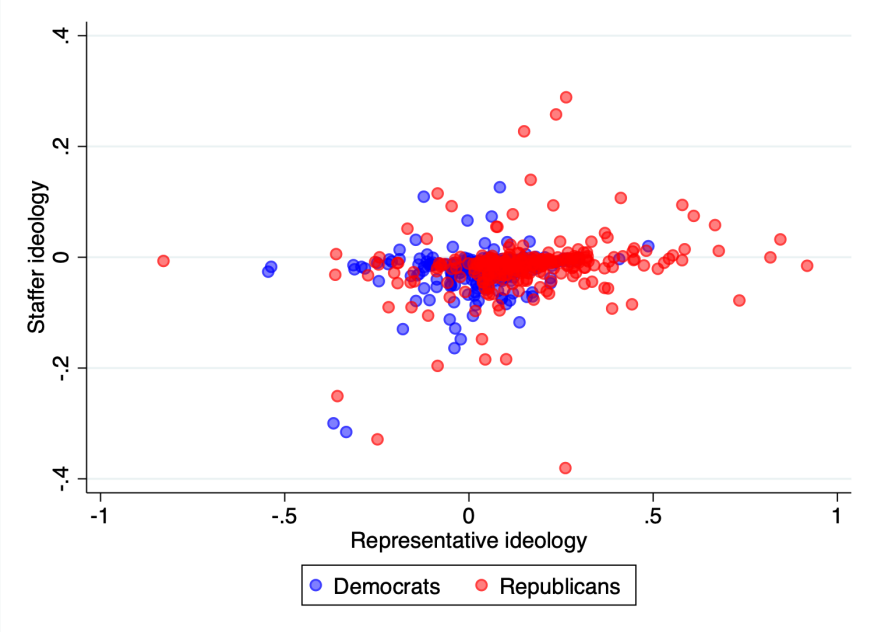
Note: This figure presents trace plots for a Bayesian estimation of the team-based mover design framework featuring a nested CES production function (Equation 10). The x-axis represents samples from a posterior drawn via MCMC. Panel A shows the office share of differences between above and below median offices in bills produced. Panel B presents the elasticity of substitution between staffers working on the same team, while Panel C presents the elasticity of substitution between the team of staffers and the Representative. Values for 4 different chains are displayed.

Figure A.10: Trace plots for ideology model, office share of differences



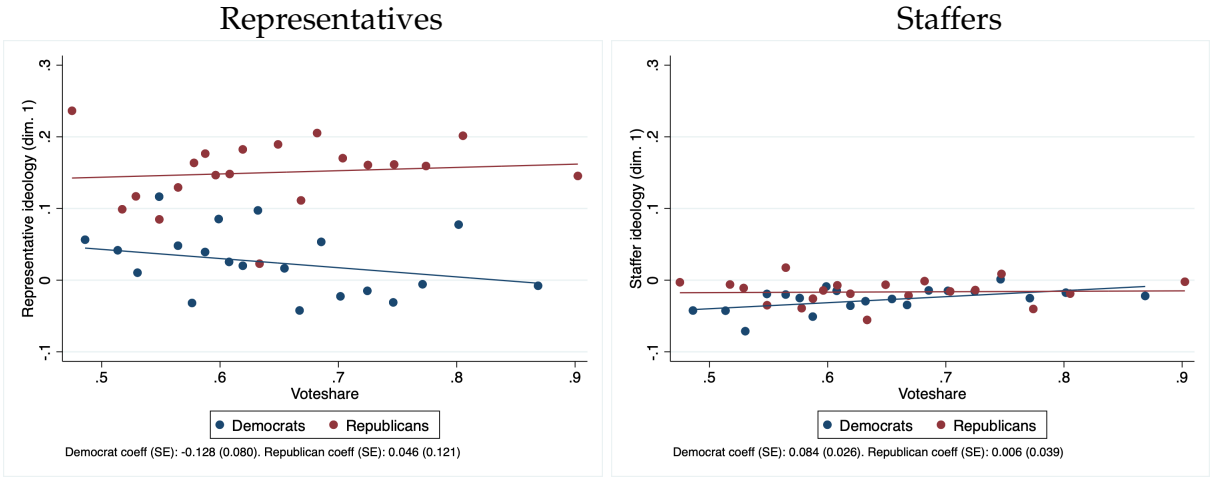
Note: This figure presents trace plots for a Bayesian estimation of the ideology of bills produced by an office (Equation 11), which is a weighted average of office (Representative) and staffer ideal points. Weights are estimated from the team-based mover design (Equation 1). The x-axis represents samples from a posterior drawn via MCMC. Both panels show the office share of differences between above and below-median offices (within party) in ideology. Panel A estimates ideology for the first DW-Nominate dimension, the conventional left-right partisan ideology. Panel B estimates ideology for the second DW-Nominate dimension, concentrated on other policy domains. Values for 4 different chains are displayed.

Figure A.11: Representative and staffer ideologies (first dimension posterior means)



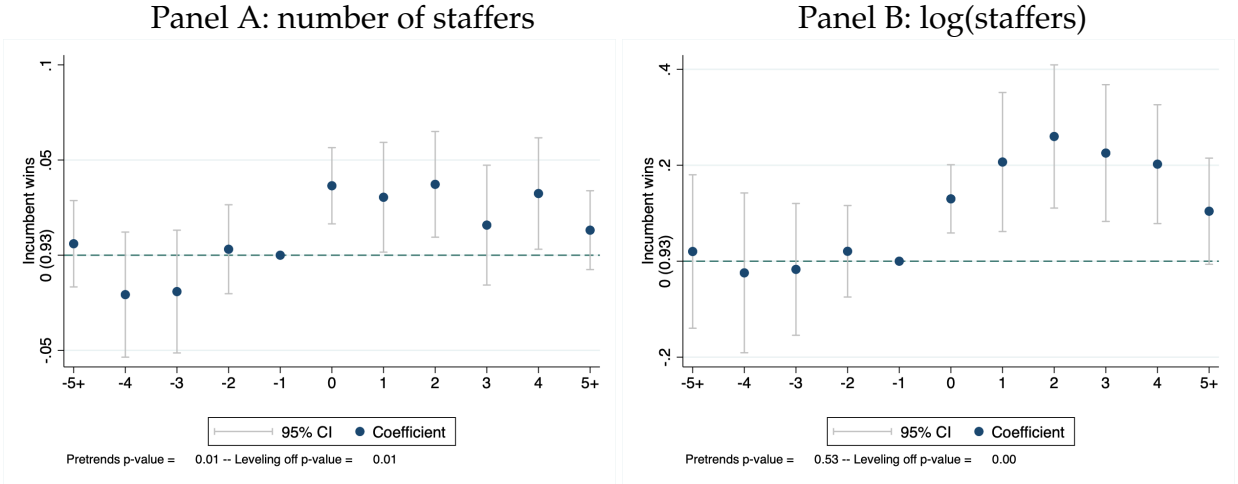
Note: This figure presents posterior means for the estimated ideology of Representatives and staffers. The staffer ideology is the (productivity-weighted) average across staffers who work for the Representative. Democrat Representatives are plotted in blue, Republicans in red.

Figure A.12: Representative and staffer ideologies by own voteshare



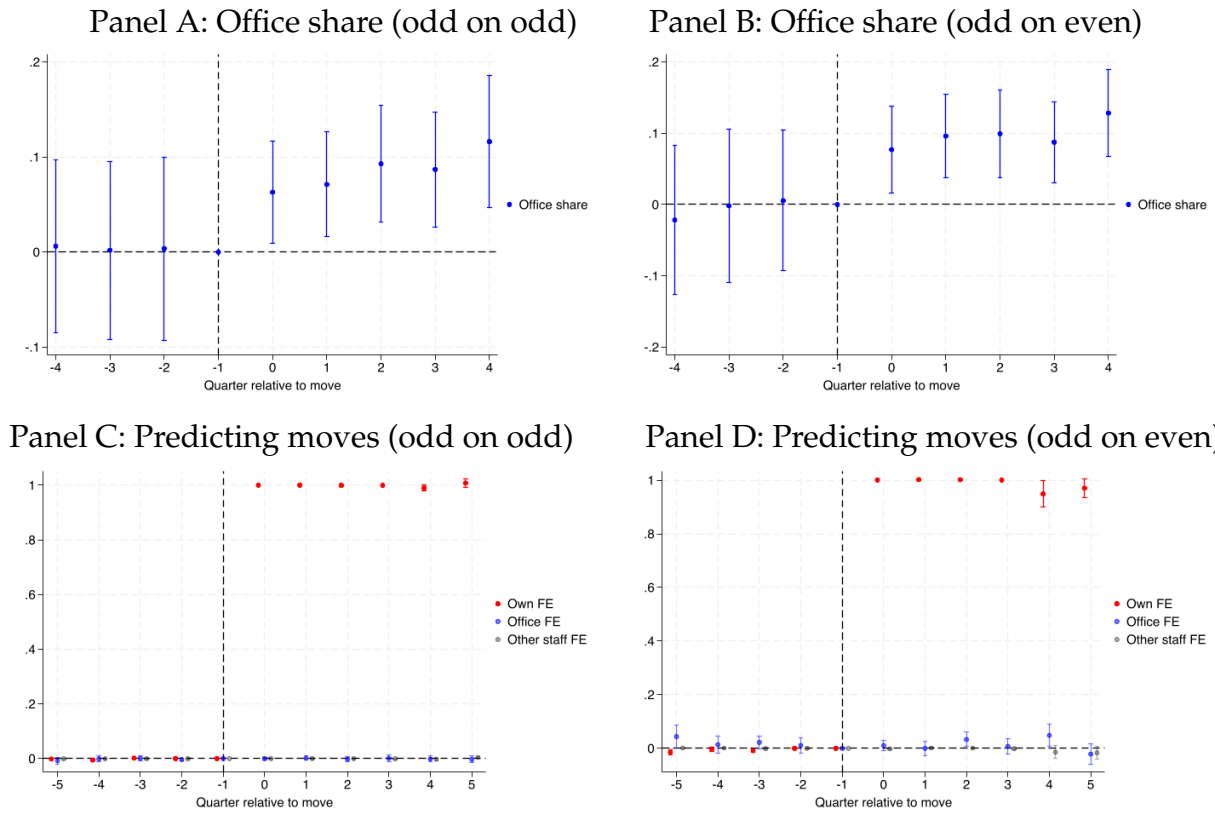
Note: This figure presents binscatter plots for posterior means of Representative (left) and staffer (right) estimated first dimension ideologies against the Representative's average voteshare. Regressions are at the Representative level, with staffer ideologies the (productivity-weighted) average across staffers who work for the Representative. Democrat are plotted in blue, Republicans in red. Standard errors for the regression coefficients presented are robust.

Figure A.13: The effect of additional staffers on the incumbency advantage



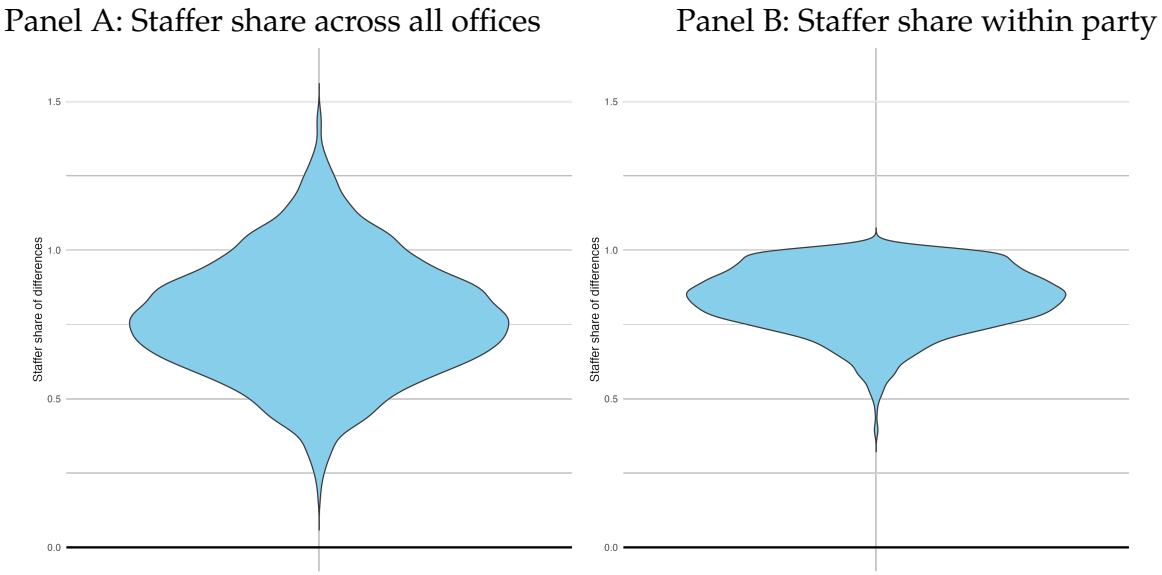
Note: This figure presents coefficients and 95% confidence intervals from a two-way fixed effect event study design. Each event is an expansion in the number of staffers that Representatives in the House were legally allowed to hire (of which there were 15 total). The regression is at the Congressional district election year level, with Senators included for the estimation of time-fixed effects. The outcome is an indicator of the incumbent winning an election. Panel A plots the coefficient for the maximum number of staffers allowed to be hired, while Panel B plots the coefficient for the log of the maximum number of staffers allowed to be hired. Standard errors are clustered at the Congressional district level.

Figure A.14: Validation of dynamic estimation by simulation



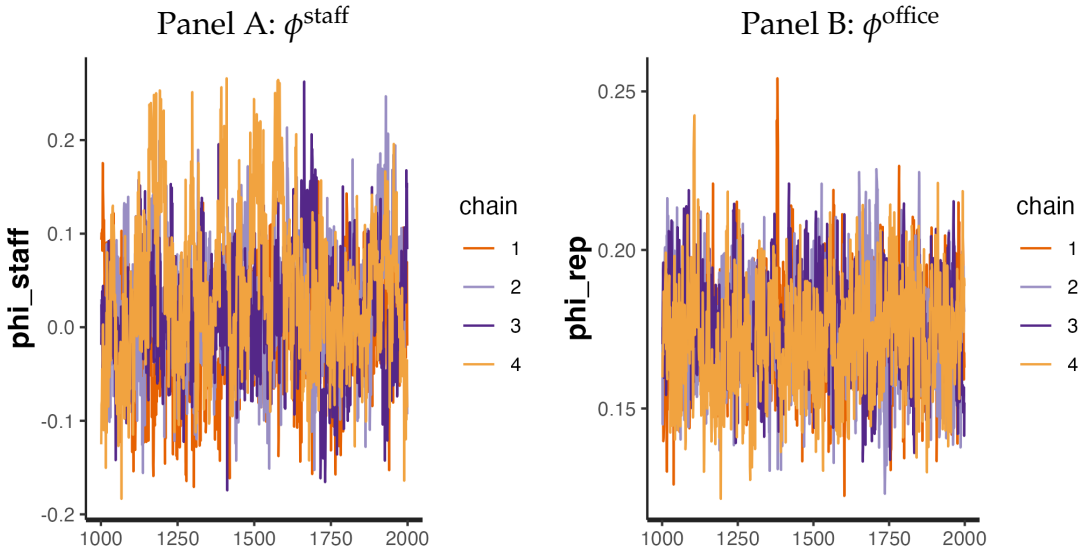
Note: This figure plots coefficients and 95% confidence intervals for two event study specifications (Equation 6 in Panels A and B, Equation 7 in Panels C and D) at the staffer-quarter level. The underlying data are simulated. The x-axis shows quarters to a staffer's first move. The outcome is bills produced. In Panels A and B, coefficients on the office share are plotted. The true values are 8 and 12% respectively. In Panels C and D, coefficients on the moving staffer's fixed effect are plotted in red, coefficients on the destination office's fixed effect are plotted in blue, and coefficients on the destination team of staffers' fixed effect are plotted in gray. Standard errors are clustered at the office level.

Figure A.15: Violin plots for ideology model, staffer share of differences on 2nd dim.



Note: This figure presents violin plots for a Bayesian estimation of the ideology of bills produced by an office (Equation 11), which is a weighted average of office (Representative) and staffer ideal points. Weights are estimated from the team-based mover design (Equation 1). Both panels show the staffer share of differences between above and below-median offices in ideology for the second DW-Nominate dimension, orthogonal to left-right ideology. Panel A shows the staffer share of differences across all offices, while Panel B shows the staffer share of differences for above and below-median offices within a party.

Figure A.16: Trace plots for experience model



Note: This figure presents trace plots for a Bayesian estimation of the team-based mover design framework featuring returns to time spent in Congress (Equation 12). The x-axis represents samples from a posterior drawn via MCMC. Panel A shows the return for staffers and Panel B shows the return for the office (Representatives). Values above 0 indicate positive returns to experience. Values for 4 different chains are displayed.