Participation, Not Adjustment: Reappraising Partisan Constraints on PAC Fundraising

Word Count: 6,500

Abstract

Businesses spend far less on campaign contributions than the law allows, a puzzle Li (2018) explains through internal constraint: access-seeking PACs must raise funds through voluntary donations from their sponsor firms' employees, and employees withhold donations when PACs give to the out-party. Li identifies this mechanism on observed donors. I revisit the claim using an original panel of 205,522 individual-PAC-cycle observations linking employment records, voter registration, and DIME contribution histories across 2012-2022 - the full eligible-employee population rather than only prior donors. Li's within-donor mechanism holds on ever-donors but does not extend to the broader sample. A specification exploiting cross-individual variation identifies a precisely estimated extensive-margin response ($\hat{\delta} = 0.035$, $p < 10^{-7}$): co-partisan employees are substantially more likely to enter donation than out-partisan ones. Mechanisms identified on behaviorally-selected samples - donors, voters, activists - often fail to extend to the populations they are meant to characterize. Population-scale claims require population-scale data.

Political scientists have long grappled with the puzzle first stated by Ansolabehere, de Figueiredo, and Snyder (2003): why is there so little money in U.S. politics? U.S. corporations spend far less on campaign contributions than the law allows, and the standard inference - that PAC contributions must generate minimal returns - sits uneasily alongside the evident political stakes. The question has animated a large subsequent literature on who gives, how much, to whom, and why: on the concentration of donations among an elite slice of the electorate (Bonica 2016; Bonica and Rosenthal 2018), on the partisan ideological extremity of individual donors relative to the broader public (Barber 2016; Hill and Huber 2017; Broockman and Malhotra 2020), and on the tangled relationships between corporate political activity, interest-group strategy, and democratic representation (Bonica 2016; Hertel-Fernandez 2018). Whether U.S. corporations' political contributions reflect optimal rent-seeking, a constrained response to internal pressures, or something else remains a first-order question for scholars of campaign finance, interest group politics, and distributive democratic representation.

Li (2018) offers a novel answer to the puzzle, grounded in internal constraint rather than external incentive: access-seeking PACs must raise funds through voluntary donations from their sponsor firms' employees, and employees withhold donations when PACs contribute to politicians of the opposite party. On Li's panel of PAC donors, this constraint operates substantially. When a PAC shifts its contributions toward one party, donors who prefer the other party reduce their probability of giving by approximately 1.1 percentage points per standard-deviation shift in PAC partisan lean (Li 2018, Table 1, Column 1). The finding has shaped subsequent debates about corporate political activity and the limits of firm fundraising (Stuckatz 2022; Li and DiSalvo 2023), providing a microfoundation for the "too little money" puzzle that does not require appealing to minimal policy returns.

Li's finding rests on a design that is simultaneously powerful and circumscribed. Her sample consists of individuals observable in FEC itemized contribution records as having donated to an access-seeking PAC. This sample is powerful because it permits tight identification of within-donor responses to changing PAC behavior - a real behavioral parameter that, as I will show, characterizes the donor sub-population correctly. It is circumscribed because it excludes the vast majority of individuals who might give to these PACs - the eligible employees who have not (yet) donated, or who have donated below the FEC itemized threshold. Li's design, like most of the empirical literature on partisan donor behavior, identifies mechanisms among those already participating. Whether the same mechanisms operate on the broader population of eligible donors - the individuals whose entry into (or abstention from) PAC giving shapes aggregate fundraising outcomes - has remained an open question.

This gap is not an oversight. It reflects a fundamental data constraint. Studying partisan preferences and PAC donation behavior at the population level requires linking two things that U.S. campaign finance data does not naturally connect: the universe of eligible donors to a PAC (who are employees of its sponsor firm) and the universe of observed donations. FEC records capture donations but observe the donor's employer only through self-report on itemized contribution records - a field Bonica (2016) shows is present only for a fraction of itemized donors and missing entirely for those who do not itemize. The Database on Ideology, Money in Politics, and Elections (DIME; Bonica 2024) has partially addressed this by standardizing and cleaning contributor metadata across cycles, but it cannot observe individuals who have never given. In countries with comprehensive administrative linkages between employment records and political contributions - the Scandinavian countries, for example - this problem is tractable by design. In the United States it has not been. Previous efforts to study partisan preferences among eligible employees have relied on sector-specific professional rosters: Bonica and co-authors have used the American Medical Association Masterfile (Bonica, Rosenthal, and Rothman 2014) and the National Provider Identifier directory (Bonica et al. 2020) to study physicians, and related work has built similar rosters for lawyers and judges (Bonica and Sen 2017). These approaches work where a comprehensive external roster exists, but they do not generalize to the general white-collar workforce that populates the corporate PAC universe. LinkedIn - a near-universal professional identity platform for white-collar workers that records employment spells with start and end dates - offers the first opportunity to construct a similar external roster for the broader corporate employee population.

This paper introduces a new empirical approach built on that opportunity. I construct a panel of 205,522 individual-PAC-cycle observations covering 840 access-seeking PACs and 63,829 individual-PAC pairs across 2012-2022, linking three data sources previously studied in isolation. Revelio Labs provides LinkedIn-derived employment records that identify, for each of Li's PAC-sponsoring firms, the universe of individuals employed at that firm in each cycle, including those who have never itemized a political contribution. L2 voter registration records provide party registration for these individuals in the 30 states plus the District of Columbia where party registration is available. DIME provides comprehensive contribution records, linked to individuals via name-and-address matching. This triple linkage allows me to observe the full population of eligible donors to each PAC, rather than only the sub-population observable as having already donated. It is the absence of this linkage that has, until now, made population-level analysis of partisan donor behavior in the corporate PAC universe infeasible.

Three findings follow from this design. First, Li's within-donor partisan response holds on her sample of ever-donors, but does *not* extend to the broader eligible-employee population. On the sample in this paper, Li's specification yields a coefficient statistically indistinguishable from zero ($\hat{\beta} = 0.006$, $p = 0.39$); restricting to pairs who have donated at least once recovers Li's magnitude ($\hat{\beta} = 0.21$). A calibrated Monte Carlo simulation rules out the explanation that this recovery is mechanical: given a small underlying parameter, selection on the ever-donor outcome cannot produce coefficients of Li's magnitude. The within-donor partisan response is a real behavioral feature of the donor sub-population - not the broader eligible-employee population from which it is drawn.

Second, a specification that Li's design structurally cannot estimate - one that identifies the alignment effect from cross-individual variation within PAC-cycle rather than from within-pair variation over time - recovers a large and precisely estimated population-level partisan response. A co-partisan eligible employee is substantially more likely than an out-partisan employee to donate in a given PAC-cycle ($\hat{\delta} = 0.035$, $p = 1.3 \times 10^{-8}$). This is the extensive margin of partisan PAC donations: the effect of alignment on the decision to participate in PAC giving at all, identifiable only on a sample that includes eligible non-donors. Li's contention that partisan preferences constrain access-seeking PAC fundraising holds at the population scale, but the binding margin is participation itself, not the within-donor decision to give in any particular cycle.

Third, this extensive-margin result is stable across cycles and robust to a range of identification controls. The coefficient is positive and individually significant in every cycle from 2012 through 2022, attenuates modestly under the most demanding state-role-seniority fixed effects specification (to $\hat{\beta} = 0.028$), and sits 13.4 standard deviations above a within-PAC permutation null. The mechanism is structural, not tied to any single cycle's political environment or to any observable correlate of partisan-aligned employee selection into specific firms.

Two methodological lessons follow. First, mechanisms identified on samples defined by the very behavior under study - donors, voters, primary participants - need not characterize the broader eligible populations from which those samples are drawn. Hill and Huber (2017), merging administrative donor records with survey data, show that donors diverge from non-donor co-partisans on a wider set of dimensions than voters diverge from non-voters - a caution specifically about the donor sub-sample that is the standard sample for studying political contributions. Bonica, Rosenthal, and Rothman (2014) and Bonica et al. (2020), studying physicians through the AMA Masterfile and the NPI directory (rather than through donor records alone), document patterns of partisan sorting that the donor-only sub-sample of physicians had obscured. Simonovits et al. (2021), linking USDA program payment records to individual voting histories, identify participation effects that survey-based studies of program beneficiaries had missed. In each case, expanding from a behavior-conditioned sample to a broader eligible population changed what the data had to say about the substantive question.

Second, where the within-participant mechanism diverges from the population-level mechanism, the gap typically runs through the extensive margin. A within-participant dose-response can appear substantively large on a behavior-conditioned sample and shrink toward zero on the broader eligible population, because the population-relevant action is concentrated in who enters that sample at all. Anzia, Jares, and Malhotra (2022) link USDA administrative records on agricultural subsidy receipt to survey responses and find that policy-feedback effects documented in much of the participant-only literature do not extend to agricultural producers; Jares and Malhotra (2025) link administrative records on individual Market Facilitation Program payments to voter-registration data and find no evidence that variation in policy benefits mobilized recipients on a mass scale. The within-donor mechanism Li measures fits this pattern: it is a real dose-response among already-participating donors but does not extend to the broader eligible-employee population. The binding constraint operates at the extensive margin, through who participates, not through the dose-response of those already inside.

The remainder of the paper proceeds as follows. Section 1 describes the data infrastructure and panel construction. Section 2 presents the empirical strategy, contrasting Li's within-pair specification with the cross-individual specification. Section 3 reports the main results: validation of the voter-registration partisanship measure (Section 3.1), the decomposition showing Li's parameter is real on her sub-population but does not generalize (Section 3.2), the extensive-margin specification and its central finding (Section 3.3), and robustness checks (Section 3.4). Section 4 concludes.

1 Data

1.1 PAC universe

I follow Li (2018) in defining access-seeking PACs as the set of corporate-sponsored PACs in OpenSecrets "business PAC" classification, totaling 5,284 PACs across the 1990-2016 cycles (though I extend the data through 2022). The substantive contribution of this paper requires linking these PACs to their sponsor firms and, in turn, to the eligible employees who could give to each PAC. I achieve this linkage in two steps: first matching each PAC to its sponsoring firm, and then matching firm employees to voter registration and donation records.

I construct the PAC-to-firm crosswalk by fuzzy matching cleaned PAC names against employer names appearing in L2 voter registration records, using a Jaro-Winkler similarity threshold (≥ 0.97 for "Tier 1" exact matches, ≥ 0.88 for "Tier 2" close matches, with shorter matches dropped to guard against over-stripping; full procedure in Appendix A.2). The procedure successfully links 2,162 of Li's 5,284 PACs to identifiable sponsor firms - a 41% match rate - distributed as 1,161 Tier 1 matches and 1,001 Tier 2 matches.

The 41% match rate reflects three structural constraints rather than crosswalk failure. First, 9.7% of Li's PACs (513 of 5,284) have missing or empty committee names in the FEC source data and cannot be matched on any name-based procedure. Second, the matching strategy is biased by design toward sponsors whose employees self-identify on LinkedIn-derived employment records: corporate PACs - where a single firm anchors the match - recover at higher rates than trade-association or membership-organization PACs, which lack a single corporate sponsor. Third, the matched set is moderately biased toward larger sponsor firms (median total disbursements 2012-2022 of \$319K versus \$155K for unmatched PACs, $p < 10^{-19}$), reflecting that larger firms have more employees visible in voter registration data. Importantly, the matched and unmatched PACs are partisan-identical: matched PACs have a mean Republican-leaning share of 0.13, statistically indistinguishable from the 0.13 mean of unmatched PACs. Whatever bias the crosswalk introduces, it is orthogonal to ideology (Appendix A.3 reports the full matched-vs-unmatched comparison).

1.2 Eligible-employee universe

The set of individuals at risk of donating to a given PAC is the set of employees of the PAC's sponsor firm. I identify these employees through Revelio Labs LinkedIn-derived employment records, accessed via the Stanford GSB Redivis workflow (Jiang 2023), which contain employment spell start and end dates for each individual at each firm in the matched PAC universe. The Revelio data extends through October 2023, which sets the upper bound analysis window: cycles ending after October 2023 (i.e., the 2024 cycle) would require imputing employment status for unobserved months and are thus excluded.

I link Revelio individuals to L2 voter registration records via name and address matching, providing each individual's state of registration and party registration (see Appendix A.1 for full documentation of the multi-tier matching procedure). This linkage yields the panel of "eligible employees" for each PAC-cycle: individuals employed at the sponsor firm during the cycle who have observable voter registration. I restrict subsequent analysis to "pure partisans" - individuals registered as Republican or Democrat - coded as $\text{Indiv\_R}_i = +0.5$ for Republicans and -0.5 for Democrats. This restriction maps directly onto Li's pure-partisan sub-sample, which she defines analogously from donation records (see Section 3.1 for measure validation).

I further match individuals to DIME contributor records via a crosswalk between L2 voter registration records and DIME contribution histories, yielding a unique DIME identifier for each matched individual. The DIME match is essential for the decomposition analysis in Section 3.2, which requires both the voter registration partisanship measure and Li's DIME-based measure on the same individuals. Across cycles 2012-2022, the DIME-matched pure-partisan sub-sample contains 205,522 pair-cycle observations: 63,829 individual-PAC pairs across 840 PACs.

1.3 Sample characteristics

I restrict the panel to cycles 2012-2022, six biennial federal election cycles. The window is constrained at both ends by data availability. The 2024 cycle is partially covered by the October 2023 Revelio snapshot and is excluded to avoid imputing employment status for unobserved months. The 2010 cycle is excluded for a different reason: the DIME contribution database used to construct individual-level donation records has incomplete coverage of pre-2012 PAC recipient identifiers, reflecting the FEC e-filing mandate that took effect in 2012. The combination of incomplete identifier coverage in 2010 and a target donor cohort that is heavily 2012-forward (only 10.3% of the DIME matches were first active in DIME on or before 2010) yields a 2010 cycle with implausibly low donation event counts (265 pair-cycles versus 1,628 in 2012 and 1,247 in 2022). Diagnostic checks confirm that the issue is specific to the DIME representation of pre-e-filing-mandate cycles rather than a true coverage gap in the underlying FEC data, but the 2010 cycle is unrecoverable without rebuilding from raw FEC bulk files. Appendix A.4 documents the diagnostic, including the retention-rate comparison across cycles and the cohort-tilt calculation. Restricting the panel to 2012-2022 yields six cycles in which both donation data and donor cohort coverage are consistent.

I construct individual-to-PAC donation records from the DIME contribution database. For each (i, j, t) triple in the pair-cycle panel, I set $\text{Give}_{ijt} = 1$ if individual i made any itemized contribution ($\geq$ \$200) to PAC j during cycle t, and 0 otherwise. This matches Li's outcome variable definition exactly. The Entry_{ijt} variable used in Column 2 of Table 4 follows Li's construction: $\text{Entry}_{ijt} = 1$ if cycle t is the first cycle in which pair (i, j) has $\text{Give}_{ijt} = 1$; $\text{Entry}_{ijt} = 0$ for cycles before the pair's first donation; and $\text{Entry}_{ijt} = NA$ for cycles after, with never-donor pairs coded $\text{Entry}_{ijt} = 0$ throughout.

In the DIME-matched pure-partisan sub-sample, 739 of 63,829 pairs (1.1%) ever produce a donation during the panel window - a rate consistent with the underlying participation rate in itemized PAC giving among eligible employees.

I construct PAC_R_{jt} - the net Republican share of PAC j's contributions in cycle t, centered at zero - from FEC PAC-to-candidate contribution records (the bulk *pas2* files), filtered to candidates only and matched to candidate party affiliation via the FEC candidate master (*cn*) files. This follows Li's replication code, which uses candidate contribution records only (excluding party committees), yielding $\text{PAC\_R}_{jt} \in [-0.5, +0.5]$. On the matched universe across 2012-2022, the distribution has mean 0.113 and standard deviation 0.243, comparable to Li's reported mean of 0.109 and standard deviation 0.261 over her 1990-2016 window.

The final analysis panel is the cross-product of eligible employees and their employer's PAC across cycles, restricted to pure partisans observable in voter registration and to the DIME-matched sub-sample for analyses requiring both partisanship measures. Table 1 reports summary statistics. The DIME-matched pure-partisan panel contains 205,522 pair-cycle observations across 63,829 unique individual-PAC pairs and 840 PACs. Mean Give_{ijt} on this sample is 0.0088 (0.88%), reflecting the sparse nature of itemized individual donations to access-seeking PACs. The panel is unbalanced because employment spells do not span the full 2012-2022 window for all individuals; the median pair contributes 3 cycles of observation.

Table 1: Summary statistics: DIME-matched pure-partisan panel, 2012–2022

Variable	N	Mean	SD	Min	Max	Li mean	Li SD
Indiv_R_{VR} (pure partisans)	2,495,218	-0.048	0.498	-0.500	0.500	—	—
Indiv_R_{DIME} (pure-partisan donors, DIME-matched)	198,040	-0.229	0.445	-0.500	0.500	0.083	0.465
PAC_R_{jt} (PAC-cycle cells, 2012-2022)	3,668	0.113	0.243	-0.500	0.500	0.109	0.261
Give_{ijt} (DIME-matched pure partisans)	205,522	0.009	0.093	0	1	0.153	—
Entry_{ijt} (DIME-matched pure partisans, at-risk)	203,363	0.004	0.066	0	1	0.072	—

DIME-matched sample counts: 840 PACs, 61,300 individuals, 63,829 pairs, 739 ever-donor pairs, 3.22 mean cycles/pair
Panel window: 2012-2022 (6 cycles); Revelio snapshot: October 2023

2 Empirical Strategy

2.1 Li's specification

The starting point is Li's (2018) difference-in-difference specification:

$$Give_{ijt} = \alpha_{ij} + \tau_{jt} + \beta\,(PAC\_R_{jt} \times Indiv\_R_i) + \varepsilon_{ijt} \tag{1}$$

where i indexes individuals, j indexes PACs, and t indexes two-year federal election cycles. The outcome Give_{ijt} is an indicator for whether individual i made any itemized contribution ($\geq$ \$200) to PAC j in cycle t. The interaction term $\text{PAC\_R}_{jt}\times$ Indiv_R_i captures alignment between the PAC's partisan giving pattern in cycle t and individual i's time-invariant partisanship. The donor-PAC pair fixed effect α_{ij} absorbs all time-invariant heterogeneity between pairs; the PAC-cycle fixed effect τ_{jt} absorbs all PAC-specific shocks in each cycle. Standard errors are clustered at the PAC level. The coefficient β is identified from within-pair over-time variation.

Li also estimates an analogous Entry specification by substituting Entry_{ijt} for Give_{ijt}, where Entry_{ijt} equals 1 in the cycle of a pair's first donation, 0 for cycles prior, and is missing for cycles after (with never-donor pairs coded 0 throughout). This identifies how alignment shifts the cycle-specific probability of first-time donation among not-yet-donated pairs.

Both specifications depend on within-pair variation for identification, which Li's donor-only sample supplies abundantly. A modified Granger causality test following Li (2018; Appendix A.2.5) finds no evidence that future alignment shifts predict current donation behavior, supporting the parallel-trends assumption on the sample I use (Appendix A.5).

2.2 The identification problem on an eligible-employee sample

Applying Li's specification to the eligible-employee sample in this paper, however, creates a new identification problem that Li did not face. In this panel, 98.9% of individual-PAC pairs never produce a donation during the 2012-2022 window. These pairs contribute $\text{Give}_{ijt} = 0$ in every cycle they appear, and the pair fixed effect absorbs this flat zero, yielding no within-pair variation to identify β against. The 1.1% of pairs that ever donate contribute all the identifying variation for Li's specification.

This is not a technical problem - the specification is well-defined and the estimates are unbiased - but it dramatically reduces what Li's specification can tell us. The partisan response Li identifies is a within-pair response among pairs that produce donation events. On this sample, that is 739 pairs. Li's specification on this data answers a question about this sub-population (Section 3.2 shows that it does), but the sub-population is a specific and behaviorally distinctive slice of the eligible-employee universe. Any statement about how partisan alignment shapes PAC fundraising at the population scale requires identifying the mechanism on a broader sample, which in turn requires a different source of variation.

2.3 Cross-individual specification

I propose an alternative specification that exploits cross-individual variation within PAC-cycle instead of within-pair variation over time:

$$Give_{ijt} = \tau_{jt} + \delta\,(PAC\_R_{jt} \times Indiv\_R_i) + \varepsilon_{ijt} \tag{2}$$

The specification is identical to Equation 1 in outcome, regressor, clustering, and sample, and differs only in the fixed-effects structure: the pair fixed effect is dropped. With only PAC-cycle fixed effects, the coefficient δ is identified from cross-individual variation within each PAC-cycle. Within a given PAC in a given cycle, this regression asks: are eligible employees whose partisanship aligns with the PAC's giving more likely to donate than eligible employees who are misaligned? The variation that drives this estimate comes from contrasts between individuals who do versus do *not* donate in a given PAC-cycle, and the vast majority of that variation - given the sparsity of donation events in the panel - is between eligible employees who have never donated and those who have.

The parameter δ therefore identifies what I call the "extensive-margin" partisan response: the effect of alignment on the decision to participate in PAC giving at all, rather than on the decision to give in any particular cycle conditional on having entered.

This specification is not available on Li's sample. Her data, by construction, consists of pairs who have already donated to an access-seeking PAC. Without never-donor pairs in the sample, there is no cross-pair contrast between donors and non-donors to identify off of; what remains is within-pair variation over time, which is what Equation 1 exploits. The richer donor universe in Li's panel makes Equation 1 well-identified but makes Equation 2 inestimable. The more sparsely populated donor universe in this panel makes Equation 1 underpowered but makes Equation 2 tractable.

2.4 Decomposition

To isolate whether the gap between Li's coefficient and the coefficient I obtain on this sample reflects a genuine population-level difference or artifactual differences in measurement, sample, or specification, I estimate Equation 1 on four progressively modified samples. The decomposition tests three potential drivers: the voter-registration partisanship measure (contrast A vs. B), expansion from the narrow overlap sample to the full DIME-matched subsample (B vs. C), and restriction to ever-donor pairs (C vs. D). Section 3.2 reports the results.

3 Results

3.1 Validation

Before the main analyses, I first verify that voter registration reliably reproduces Li's DIME-based measure of individual partisanship. Li's Indiv_R_i is the net share of a donor's lifetime direct contributions to Republican candidates and party committees, centered at zero, and she restricts her main analyses to "pure partisans" - individuals who have given only to one party ($\text{Indiv\_R}_i \in -0.5, +0.5$). My measure is a voter-registration analogue: +0.5 for individuals registered as Republican, -0.5 for individuals registered as Democrat, with independents and non-partisan registrants excluded.

Table 2: Validation: Voter registration vs. DIME-inferred partisanship

	DIME party			
Voter registration	Democrat	Republican	Total	Agreement
Democrat	142,348	1,065	143,413	99.26%
Republican	1,982	52,645	54,627	96.37%
Total	144,330	53,710	198,040	98.46%

Notes: Cells are pair-cycle observations in the 2012–2022 DIME-matched pure-partisan panel where both voter-registration party and DIME-inferred party (Li's $Indiv\_R_i, \in \{-0.5, +0.5\}$) are defined. Overlap sample: 198,040 pair-cycles across 59,012 individuals. Agreement column is within-row: $N_{DD}/(N_{DD} + N_{DR})$ for Democrats, $N_{RR}/(N_{RD} + N_{RR})$ for Republicans, overall $(N_{DD} + N_{RR})/N$.

Table 2 reports agreement between the two measures on the overlap sample - individuals observable in both the voter-registration data and DIME. Of pair-cycle observations where both measures are defined, 98.46% assign the individual to the same party. The rare disagreements (1.54%) are concentrated on individuals with very few DIME observations, where a single contribution flips the classification. Treating voter registration as a substitute for DIME's donation-based classification sacrifices almost nothing in terms of partisan identification, and gains the ability to classify the 84% of the sample that DIME does not observe. The remaining analyses use the voter-registration measure throughout.

3.2 Evaluating Li's original finding

Li reports a within-donor partisan response of $\hat{\beta} = 0.0853$ on pure partisans (her Table 1, Column 1). I establish two things about this estimate. First, it is a real behavioral parameter on her sub-population - eligible employees who have selected into donation by at some point donating to an access-seeking PAC. Second, this same behavioral response does not characterize the broader eligible-employee population from which her sample is drawn.

Table 3 reports four rows, each estimating Li's Equation 1 with progressive sample modifications. Row A uses the sample with the DIME-based partisanship measure: $\hat{\beta} = 0.0050$ (SE = 0.0068). Row B repeats with the voter-registration measure: $\hat{\beta} = 0.0060$, statistically indistinguishable from Row A. Row C expands the sample to all DIME-matched individuals classifiable via registration: $\hat{\beta} = 0.0059$, again indistinguishable. Neither the measurement swap (moving from Row A to Row B) nor the sample expansion (moving from Row B to Row C) meaningfully changes the coefficient. Li's specification, applied to the eligible-employee population, identifies essentially no within-donor partisan response.

Table 3: Endogeneity decomposition: Li's specification under progressive sample modifications

Row	Sample	Measure	$\hat{\beta}$	SE	N	% of Li
A	Overlap (N=198,040)	DIME pure-partisan	0.0050	0.0068	198,040	5.8%
B	Overlap (N=198,040)	Voter registration	0.0060	0.0069	198,040	7.0%
C	Full DIME-matched (N=205,522)	Voter registration	0.0059	0.0068	205,522	6.9%
D	Overlap, donor-conditioned (N=3,656)	Voter registration	0.2076	0.3390	3,656	243.4%
D'	Overlap, donor-conditioned (N=3,656)	DIME pure-partisan	0.1393	0.3199	3,656	163.3%
Li	Li (2018)	DIME pure-partisan	0.0853	0.0080	2,600,000	100.0%
Endogeneity (A−B)			−0.0011			
Sample expansion (B−C)			+0.0001			
Donor conditioning (C−D)			−0.2017			
Measure check within donors (D−D')			+0.0683			

Row D, however, restricts the Row B sample to ever-donor pairs - those in which the donor contributed to the PAC at least once during the panel window. N falls from 198,040 to 3,656 across 717 pairs (the overlap-sample count, slightly fewer than the 739 ever-donor pairs on the full DIME-matched panel, reflecting the narrower sample on which both partisanship measures are defined). The coefficient rises to $\hat{\beta} = 0.208$ (SE = 0.339) — more than twice Li's benchmark and well within the range consistent with her estimate, though imprecisely estimated on the 717-pair sub-sample. The entire gap between the broader-population estimates ($\hat{\beta} \approx 0.006$) and Li's collapses into this single sample restriction; the Monte Carlo result discussed below (and Appendix A.6) confirms that the magnitude is genuinely large rather than mechanically inflated by the small sub-sample. The regression specification, the partisanship measure, and the cycle window are all unchanged across rows; only which pairs the regression sees changes from C to D.

What this isolates is something real about the ever-donor sub-population. Conditional on having entered PAC giving, eligible employees are substantially more responsive to within-pair shifts in PAC alignment than the broader eligible-employee population would suggest. The contrast is not subtle: the within-donor response on ever-donors is roughly 35× the response on the full DIME-matched sample, and it survives the same fixed-effects discipline Li applies. Whatever distinguishes employees who have crossed into PAC giving from those who have not, it includes a behavioral sensitivity to partisan alignment in the within-donor margin that the broader population does not exhibit.

A natural concern is that the recovered coefficient on ever-donors is mechanical rather than behavioral. Restricting to pairs with at least one observed donation selects pairs with multiple Give events across cycles, providing within-pair variation that the never-donor majority does not provide. The within-pair coefficient on ever-donors might therefore be large because that is exactly where the within-pair variation mechanically lives, not because the underlying response is genuinely larger.

A calibrated Monte Carlo simulation rules this out (Appendix A.6 documents the full procedure). I simulate pair-cycle panels with a known small within-donor parameter ($\delta = 0.02$) and vary the baseline participation rate p_0, matching the cycle count, PAC count, and alignment distribution to the panel. For each p_0, I estimate Li's specification on the simulated ever-donor sub-sample across 100 draws. At $p_0 = 0.15$ - Li's reported participation rate - the simulated ever-donor coefficient has a mean of 0.0316 and a standard deviation of 0.0067. Li's 0.0853 is 8.0 standard deviations above this mean; no simulated draw produces a coefficient as large. A small underlying behavioral parameter, combined with selection on the ever-donor outcome, cannot mechanically inflate the coefficient to Li's magnitude. The within-donor response on the ever-donor sub-population is genuinely large, not amplified by sample restriction.

Two findings follow. First, on eligible employees who have already entered PAC giving, partisan alignment substantially shifts within-donor decisions to give in any particular cycle. The behavioral mechanism is real; Li identified it correctly. Second, the same mechanism *does* not characterize the broader eligible-employee population. Where Li's specification yields $\hat{\beta} = 0.21$ on ever-donors, it yields $\hat{\beta} = 0.006$ on the eligible-employee sample as a whole. The within-donor response Li measured is specific to the donor sub-population - a feature of how people who have already decided to give respond to alignment, not a feature of how partisan alignment shapes the broader decision to participate at all. Section 3.3 identifies what does shape the broader decision.

3.3 The extensive-margin mechanism

Section 3.2 identified a behavioral feature of the donor sub-population: conditional on having entered PAC giving, eligible employees are substantially more responsive to within-pair alignment shifts than the broader eligible-employee population. The remaining question is what shapes the prior decision to enter that sub-population - what partisan alignment does at the participation margin where Li's design cannot see.

Table 4 reports three specifications. Columns 1 and 2 harmonize directly with Li's Table 1 (Give) and Table 2 (Entry): identical regression equation, identical fixed-effects structure (donor-PAC pair plus PAC-cycle), identical clustering at PAC, differing from Li only in the sample to which they are applied. Column 3 presents Equation 2, which drops the pair fixed effect and identifies the alignment effect from cross-individual variation within each PAC-cycle - a specification Li could not estimate on her donor-conditioned sample because it lacks the never-donor observations that contribute the cross-pair contrast.

On the eligible-employee panel, Li's Give specification (Column 1) yields $\hat{\beta} = 0.0059$ (SE $= 0.0068$, $p = 0.39$): positive, in the direction Li reports, but not statistically distinguishable from zero. The Entry specification (Column 2) yields $\hat{\beta} \approx 0$ (SE $= 0.0045$, $p = 0.998$) - precisely zero. These specifications do what Li's specifications do: they identify a partisan

Table 4: Main results: Within-pair and cross-individual specifications on the eligible-employee panel

	(1) Give (Li Tab 1)	(2) Entry (Li Tab 2)	(3) Equation 2 (novel)
Alignment	0.0059	0.0000	0.0352***
	(0.0068)	(0.0045)	(0.0062)
t	0.87	0.00	5.69
Outcome	Give_{ijt}	Entry_{ijt}	Give_{ijt}
Fixed effects	pair + PAC×cycle	pair + PAC×cycle	PAC×cycle (no pair FE)
SE cluster	PAC	PAC	PAC
N obs	188,205	185,695	204,855
N unique pairs	63,829	63,829	63,829
N PACs	840	840	840
Li benchmark (pure)	0.0853 (0.0080)	0.0400 (0.0046)	—
Ratio to Li	0.069	0.000	—

response from within-pair variation over time. But on this sample, they have little within-pair variation to work with. Only 739 of 63,829 individual-PAC pairs in the DIME-matched panel ever produce a donation during the panel window. The remaining 98.9% of pairs contribute a flat zero to the within-pair contrast, and the pair fixed effects absorb the cross-pair variation that would otherwise carry the partisan signal. The within-donor mechanism documented in Section 3.2 - large on the ever-donor sub-population - is not detectable on the broader eligible-employee population using Li's specifications.

Column 3 reaches that signal that Li's design absorbs. Replacing the pair fixed effect with PAC-cycle fixed effects alone, the alignment coefficient is $\hat{\delta} = 0.0352$ (SE = 0.0062, $p = 1.3 \times 10^{-8}$). The only difference between Columns 1 and 3 is which source of variation identifies the effect. Column 1 identifies from within-pair timing among the $\sim 1\%$ of pairs that ever donate. Column 3, instead, identifies from cross-individual differences within each PAC-cycle: within a given PAC in a given cycle, are eligible employees whose partisanship aligns with the PAC's giving more likely to donate than employees whose partisanship misaligns? The answer is yes, substantially, and the coefficient is precisely estimated because nearly all the identifying variation comes from contrasts between eligible employees who do versus do not cross the threshold of participation at all.

This is the extensive margin of partisan PAC donations, and it is where the population-level partisan response operates. In a given PAC-cycle, a co-partisan eligible employee is substantially more likely than an out-partisan eligible employee to enter donation. Li's contention - that partisan preferences constrain access-seeking PAC fundraising - holds at the population scale, but the binding margin is participation itself, not the within-donor decision to give in any particular cycle. Her Equation 1, applied to a sample composed almost entirely of already-participating donors, identifies how alignment shifts within-donor giving among that sub-population - and on her sample, it does. Equation 2 in this paper, applied to a sample that includes the eligible employees who are not yet donors, identifies how alignment shifts the prior decision to participate at all - and on the sample here, the response is precise and substantively meaningful.

The pooled coefficient $\hat{\delta} = 0.0352$ implies that a one-unit increase in the alignment interaction raises the per-cycle probability of donation by 3.5 percentage points; against a baseline rate of 0.88% in the DIME-matched sample, this is a substantial multiplier on participation. A one-standard-deviation shift in PAC_R_{jt} (0.24) for a pure-partisan individual ($\text{Indiv\_R} = \pm 0.5$) corresponds to an alignment change of about 0.12 and raises the per-cycle donation probability by about 0.4 percentage points - roughly half the baseline rate. Aggregated across the 63,829 pairs and 6 cycles in the panel, the partisan response at the extensive margin shapes a population-level fundraising pattern of substantial magnitude: approximately \$108,000 in suppressed donation revenue per cycle across the matched PACs (Appendix A.8; Appendix A.8 also reports an upper-bound extrapolation to Li's full 5,284-PAC universe, which is sensitive to assumptions about the distribution of unmatched-PAC sizes).

Figure 1 plots the Column 3 coefficient by individual cycle. The alignment effect is positive and individually significant in every cycle from 2012 through 2022, with point estimates ranging from $\hat{\delta} = 0.024$ (2022) to $\hat{\delta} = 0.045$ (2018). The coefficient is broadly stable across cycles, with a modest rise through 2018, a plateau in 2020, and an attenuation in 2022

- consistent with the post-January 6 corporate-PAC freezes that reduced employee giving to many of the PACs in the sample. The cycle-by-cycle stability of the extensive-margin response is itself evidence that the mechanism is a structural feature of partisan-aligned employee participation, not an artifact of any single cycle's political environment.

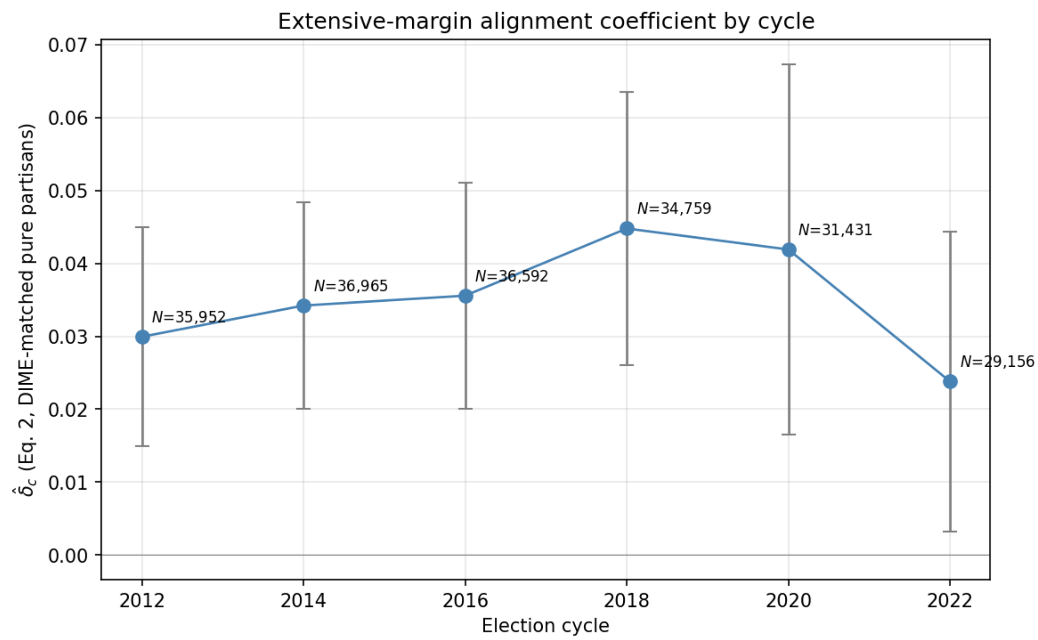

DIME-matched pure-partisan sample, 2012–2022. Equation 2 estimated separately for each cycle. Whiskers show 95% confidence intervals (PAC-clustered).

Figure 1: Extensive-margin alignment effect by cycle, 2012–2022
Note: Point estimates of $\hat{\delta}_c$ from Equation 2 estimated separately by cycle, with 95% confidence intervals. Standard errors clustered at the PAC level.

The contrast between Columns 1 and 3 of Table 4 is the paper's central empirical finding. The within-donor mechanism Li identified operates on a behaviorally distinctive subpopulation (Section 3.2) but does not extend to the broader eligible-employee sample, where it is statistically indistinguishable from zero. The extensive-margin mechanism - identified from cross-individual variation that Li's design absorbs - is large, precisely estimated, and stable across cycles. The partisan constraint Li theorized operates at the population scale through participation: through who enters PAC giving among eligible employees, not through how those who have already entered respond to alignment shifts cycle by cycle.

3.4 Robustness

The Column 3 result is robust across a range of identification choices (Table 5). Restricting the sample to Tier 1 (exact-name) crosswalk matches yields $\hat{\delta} = 0.0345$ ($p < 10^{-5}$), essentially unchanged from the baseline. Adding individual state fixed effects to absorb any state-level confounders that might correlate with both alignment and donation propensity yields $\hat{\delta} = 0.0353$. Adding state and Revelio role-category fixed effects - to absorb selection of partisan-aligned employees into specific occupations within PAC-sponsoring firms - yields $\hat{\delta} = 0.0321$, an attenuation of approximately 9%. The most demanding specification adds state, role, and seniority effects, yielding $\hat{\delta} = 0.0276$, a 22% attenuation from the baseline.

Table 5: Robustness: Extensive-margin specification with sorting controls

Specification	FE	$\hat{\delta}$	SE	t	N obs
(1) Baseline	PAC×cycle	0.0352***	(0.0062)	5.69	204,855
(2) Tier 1 restriction	PAC×cycle	0.0345***	(0.0072)	4.79	158,516
(3) + State FE	PAC×cycle + state	0.0353***	(0.0061)	5.81	204,855
(4) + State + Role FE (k50)	PAC×cycle + state + role	0.0321***	(0.0059)	5.45	204,855
(5) + State + Role + Seniority FE	PAC×cycle + state + role + seniority	0.0276***	(0.0055)	5.04	204,855

The seniority attenuation is consistent with some of the alignment-giving covariance operating through partisan-aligned selection of employees into senior roles within PAC-sponsoring firms - a mechanism that the seniority fixed effects absorb but that one would otherwise interpret as part of the partisan response. The coefficient remains highly significant ($p = 5 \times 10^{-7}$) and substantively large under the most demanding specification, but the attenuation is a real feature of the data, not noise. The substantive interpretation of the result accommodates either reading: the partisan participation response is large in either case, and the question of how much operates through occupation-specific selection versus direct alignment-driven entry is a question about mechanism that the specifications in this paper cannot fully separate.

A placebo test confirms that the result is not a chance artifact of the regressor structure. I shuffle voter registration within each PAC and re-estimate the specification 500 times; the resulting null distribution is centered at zero with a standard deviation of 0.0026 (Appendix A.7). The observed coefficient of 0.0352 is 13.4 standard deviations above the placebo mean, and not a single placebo draw produces a coefficient as large as the observed value. Whatever the Column 3 coefficient is measuring, it is not noise in the alignment construction.

4 Conclusion

This paper has reappraised Li (2018), who argued that access-seeking PACs are constrained from extreme partisan giving by the partisan preferences of the employees who fund them. I have offered three findings, each enabled by a panel that observes the eligible-employee population rather than only the donor sub-population.

First, Li's within-donor partisan response is real on her sample of ever-donors. The same regression specification, applied to the broader eligible-employee sample that her question is ultimately about, yields essentially no within-donor response. The gap between her coefficient and the one in this paper is not driven by the partisanship measure or the panel window; it is driven entirely by the sample restriction to ever-donors. A calibrated simulation confirms that this is not a mechanical artifact: even granting Li's sample construction, the magnitude she identifies cannot be produced by selection on outcome alone given a small underlying parameter. Whatever distinguishes the donor sub-population from the broader eligible-employee population, it includes a behavioral sensitivity to within-pair alignment shifts that the broader population does not exhibit.

Second, the population-level partisan mechanism Li theorized does operate - but at a different margin than her design could measure. A specification that exploits cross-individual variation within PAC-cycle, rather than within-pair variation over time, identifies a precisely-estimated extensive-margin response. Co-partisan eligible employees are substantially more likely than out-partisan ones to enter donation in any given PAC-cycle. The constraint on access-seeking PAC fundraising is binding at the population scale, but it binds through who participates, not through the dose-response of those already inside.

These findings have implications for the substantive question Li poses. The "too little money in politics" puzzle is, at its core, a population-level question about why interest group fundraising falls short of what theory predicts. Li's answer - partisan preferences of employees constrain corporate PAC fundraising - survives the reappraisal but is sharper than her data could show. The constraint is not primarily that committed donors withhold from misaligned PACs; it is that misaligned would-be donors do not enter in the first place. This distinction matters for thinking about the fundraising frontier of corporate PACs. A PAC's existing donor base is, by selection, a partisan-aligned subset of its eligible employee population; the dose-response Li identifies operates only weakly within that subset. The first-order question for a PAC seeking to expand fundraising is the cross-individual question of who among the broader eligible workforce will enter - and the partisan composition of that workforce, relative to the PAC's giving pattern, sets the population-scale ceiling on participation. Designs that condition on observed donors will tend to understate the constraint, because they cannot observe the participation margin at which it primarily binds.

These findings also speak to a methodological pattern that extends beyond Li's specific question. The limitation I document is shared by any design that identifies a partisan or behavioral mechanism on a sample defined by the very behavior under study - donors, voters, primary participants. The within-participant dose-response that these designs measure can be real, and substantively important on its own, but it cannot be assumed to characterize the broader eligible population. Where the data permit the comparison, the population-relevant action often turns out to be at the extensive margin - at who enters the participating sample at all - rather than in how those already inside respond to the dose. Population-level claims, then, should be tested on population-level samples whenever the data permits.

The data infrastructure in this paper carries limitations worth acknowledging. The eligible-employee panel is constructed through probabilistic record linkages: Revelio's identification of who works where rests on LinkedIn-derived employment spells, which under-represent workers without LinkedIn profiles and may misdate employment transitions. To the extent LinkedIn coverage correlates with partisanship, the cross-individual estimate could be biased; the within-PAC permutation placebo (Appendix A.7) rules out mechanical artifacts of the panel structure but cannot rule out coverage-correlated selection. The L2 voter registration link rests on name-and-address matching, which fails for individuals who have moved or changed name. Moreover, the voter registration data is a snapshot rather than a longitudinal record, so individuals who changed party registration during the panel window are observed at their most recent party. This raises a bounded reverse-causality concern: donating to a co-partisan PAC could plausibly cause a partisan registration update, which would inflate the observed alignment-donation correlation. The cycle-by-cycle stability of the extensive-margin coefficient (Figure 1) limits this concern: a registration-update-driven mechanism would concentrate the effect in cycles immediately following major partisan realignments, which the data does not show. I restrict to the 30 states plus DC where party registration is available, excluding California-equivalent states where partisanship must be inferred from primary voting history. Each of these decisions trades coverage for measurement precision. While these limitations are real, they should not foreclose the kind of work this paper attempts. A noisy view of the broader population is, for many questions, more informative than a clean view of a sub-population.

References

Ansolabehere, Stephen, John M. de Figueiredo, and James M. Snyder Jr. 2003. "Why Is There so Little Money in U.S. Politics?" *Journal of Economic Perspectives* 17 (1): 105–130.

Anzia, Sarah F., Jake Alton Jares, and Neil Malhotra. 2022. "Does Receiving Government Assistance Shape Political Attitudes? Evidence from Agricultural Producers." *American Political Science Review* 116 (4): 1389–1406.

Barber, Michael J. 2016. "Ideological Donors, Contribution Limits, and the Polarization of American Legislatures." *Journal of Politics* 78 (1): 296–310.

Bonica, Adam. 2016. "Avenues of Influence: On the Political Expenditures of Corporations and Their Directors and Executives." *Business and Politics* 18 (4): 367–394.

Bonica, Adam. 2024. *Database on Ideology, Money in Politics, and Elections: Public Version 4.0* [Computer file]. Stanford, CA: Stanford University Libraries. `https://data.stanford.edu/dime`.

Bonica, Adam, and Howard Rosenthal. 2018. "Increasing Inequality in Wealth and the Political Expenditures of Billionaires." Working paper, SSRN. `https://ssrn.com/abstract=2668780`.

Bonica, Adam, Howard Rosenthal, Kristy Blackwood, and David J. Rothman. 2020. "Ideological Sorting of Physicians in Both Geography and the Workplace." *Journal of Health Politics, Policy and Law* 45 (6): 1023–1057.

Bonica, Adam, Howard Rosenthal, and David J. Rothman. 2014. "The Political Polarization of Physicians in the United States: An Analysis of Campaign Contributions to Federal Elections, 1991 Through 2012." *JAMA Internal Medicine* 174 (8): 1308–1317.

Bonica, Adam, and Maya Sen. 2017. "The Politics of Selecting the Bench from the Bar:

The Legal Profession and Partisan Incentives to Politicize the Judiciary." *Journal of Law and Economics* 60 (4): 559–595.

Broockman, David E., and Neil Malhotra. 2020. "What Do Partisan Donors Want?" *Public Opinion Quarterly* 84 (1): 104–118.

Hertel-Fernandez, Alexander. 2018. *Politics at Work: How Companies Turn Their Workers into Lobbyists*. New York: Oxford University Press.

Hill, Seth J., and Gregory A. Huber. 2017. "Representativeness and Motivations of the Contemporary Donorate: Results from Merged Survey and Administrative Records." *Political Behavior* 39 (1): 3–29.

Jares, Jake Alton, and Neil Malhotra. 2025. "Policy Impact and Voter Mobilization: Evidence from Farmers' Trade War Experiences." *American Political Science Review* 119 (2): 847–869.

Jiang, Mason. 2023. *Revelio Workforce Example Project* [Workflow]. Redivis. DOI: 10.71778/V2DW-7A53. `https://stanfordgsb.redivis.com/workflows/hhb4-9x8ayakn6`.

Li, Zhao. 2018. "How Internal Constraints Shape Interest Group Activities: Evidence from Access-Seeking PACs." *American Political Science Review* 112 (4): 792–808.

Li, Zhao, and Richard W. DiSalvo. 2023. "Can Stakeholders Mobilize Businesses for the Protection of Democracy? Evidence from the U.S. Capitol Insurrection." *American Political Science Review* 117 (3): 1130–1136.

Simonovits, Gábor, Neil Malhotra, Raymond Ye Lee, and Andrew Healy. 2021. "The Effect of Distributive Politics on Electoral Participation: Evidence from 70 Million Agricultural Payments." *Political Behavior* 43 (2): 737–750.

Stuckatz, Jan. 2022. "How the Workplace Affects Employee Political Contributions." *American Political Science Review* 116 (1): 54–69.

A Online Appendix

A.1 Revelio-L2 Voter Registration Merge

Linking Revelio employment records to L2 voter registration profiles requires resolving two sources of name variation: differences between how individuals list themselves on LinkedIn versus how they appear in official voter rolls (nicknames, name truncations, middle-name usage), and geographic ambiguity when multiple individuals share a name in a large state. I address both through a multi-tier iterative procedure that proceeds from most to least restrictive matching criteria, retaining only unambiguous one-to-one links at each stage.

Pre-processing. Before matching, I restrict the Revelio pool to profiles with a non-missing first name, last name, and resolvable geographic location. I standardize each self-reported location string to a Census geographic identifier using the 2024 Census Gazetteer files, falling back sequentially from Census place to county to Core-Based Statistical Area (CBSA) if no place-level match is found. L2 residential addresses are standardized to the same hierarchy. Matching is blocked on state throughout: no record is ever matched across state lines.

Matching tiers. The procedure runs 108 rounds organized into four tiers of declining name-matching stringency. Within each tier, rounds proceed from most to least restrictive: matches requiring both birth year agreement (± 10 years) and middle initial consistency are attempted first, followed by birth year alone, middle initial alone, and finally name and geography without additional constraints. A profile is carried forward to lower-stringency rounds only if it has not been matched in any prior round.

Tier 1 requires exact agreement on both first and last name. This tier runs 36 rounds spanning all combinations of birth year and middle initial constraints across three geographic levels (Census place, county, CBSA).

Tier 2 truncates the first name to its first four characters while requiring an exact last name, recovering individuals who report a nickname or informal first name on LinkedIn (e.g., "Alex" vs. "Alexander"). The same 36-round structure is applied.

Tier 3 truncates both first and last name to four characters, accommodating minor spelling variation and data entry differences. This tier runs 36 rounds with the same constraint structure.

Tier 4 applies nickname normalization - mapping common informal names to their canonical legal forms using two independent nickname libraries - before running the same exact-name rounds as Tier 1. Two passes are run using different underlying libraries, with the higher-confidence match preferred when both recover a candidate.

Throughout all tiers and rounds, a profile is retained in the final crosswalk only if it maps to *exactly one* L2 voter record under the applicable criteria. Profiles that generate two or more candidate matches under any round are excluded entirely, trading recall for precision.

Results. Table A.1 reports the count of matched individuals at each stage of the pipeline. The 4-tier procedure produces 44,124,662 deduplicated individual matches in the production crosswalk used for analysis (Panel A), with exact-name matches accounting for the large majority. A parallel audit pipeline preserves the per-tier labels and restricts to the 2012–2022 analytical sample, yielding 680,689 tier-labeled matches distributed as Tier 1 (487,638), Tier 2 (73,847), Tier 3 (52,022), and Tier 4 (67,182); Panel B reports this breakdown. Of the 44,124,662 production matches, 14.4% are linked to a DIME contributor record via a separate name-and-address crosswalk (Bonica 2024), which inherits the L2 voter ID as a common key.

Restricting to the 2012-2022 analytical window - individuals employed at a matched PAC-sponsoring firm during at least one observed cycle and registered as a partisan in a state that records party affiliation - yields a broad eligible-employee panel of 380,153 unique individuals. Of these, 61,300 (16.1%) are matched to a DIME contributor record, forming the DIME-matched sub-sample used in the replication and validation analyses.

The tier distribution within the analytical sample reflects the heavy concentration of

Table A.1: Revelio–L2 merge pipeline

Stage	N	Share
Panel A: Revelio → L2 linkage		
Raw Revelio profiles (pre-match pool)	183,994,482	–
Eligible for matching[a]	135,949,177	73.9%
Total unique matched	44,124,662	32.5%
Exact name match	38,883,018	88.1%
Nickname match (DFS library)	4,598,599	10.4%
Nickname match (basic library)	643,045	1.5%
Panel B: Tiered name-matching breakdown		
Tier 1 — exact first & last name	487,638	71.6%
Tier 2 — first 4 + exact last	73,847	10.8%
Tier 3 — first 4 + last 4	52,022	7.6%
Tier 4 — nickname normalization	67,182	9.9%
Total (high-stringency tiered subset)[b]	680,689	100.0%
Panel C: DIME campaign-finance linkage		
Revelio–L2 matched pool linked to DIME	6,344,055	14.4%
Panel D: 2012–2022 analytical panel		
Unique individuals in panel[c]	380,153	–
Of which: DIME-matched sub-sample	61,300	16.1%

Notes:

[a] Eligible = non-null, non-empty first name, last name, and geographic match level (`place/county/cbsa`).

[b] The tiered file is a narrower, high-stringency subset used for audit of name-match quality; the production crosswalk in Panel A (44,124,662 matches) feeds all downstream analysis. The Tier 1–4 labels in Panel B are not available on the production crosswalk.

[c] Individuals employed at a crosswalk-matched PAC sponsor firm during at least one 2012-2022 cycle and registered as a pure partisan (Democrat or Republican) in a party-registration state.

recoverable matches at the most stringent name criterion. Within the 2012-2022 analytical sample, 487,638 of 680,689 tier-labeled matches are recovered via Tier 1 exact-name matching, with smaller contributions from Tiers 2–4. Table A.2 reports the full tier breakdown.

Table A.2: Crosswalk tier distribution in the 2012–2022 analytical sample

Tier	Name-match criterion	Pair-cycles	Share
Tier 1	Jaro–Winkler ≥ 0.97 (exact-name)	158,858	77.29%
Tier 2	Jaro–Winkler $\in [0.88, 0.97)$ (close-name)	46,664	22.71%
Total		205,522	100.00%

Notes: Pair-cycle observations in the 2012–2022 DIME-matched pure-partisan analytical panel, broken down by the DIME–PAC name-matching tier at which the PAC was linked to its sponsor firm. Tier 1 covers 470 unique PACs; Tier 2 covers 370 unique PACs; total 840 PACs. These tiers pertain to the DIME–PAC (name-matching) linkage, not the Revelio–L2 (individual-matching) tiers reported in Table A.1.

Validation. The main results are robust to restricting the analytical sample to Tier 1 exact-name matches only. The extensive-margin alignment coefficient on the Tier 1 exact-name sub-sample (487,638 matches feeding a restricted analytical panel) is $\hat{\delta} = 0.0345$ ($p < 10^{-5}$), compared to the full-sample baseline estimate of 0.0352. The near-equivalence indicates that lower-stringency tiers introduce at most modest attenuation noise rather than directional bias, and that the headline finding does not depend on the inclusion of less-confident matches.

DIME campaign finance linkage. Individual PAC donation records come from the Database on Ideology, Money in Politics, and Elections (DIME; Bonica 2024). I link DIME contributor records to the Revelio-L2 matched pool using a pre-built probabilistic crosswalk constructed from name and residential address fields common to both DIME and L2 (Bonica 2024). Because both the Revelio-L2 crosswalk and the DIME-L2 crosswalk share the L2 voter ID as a common key, the join is deterministic conditional on both crosswalks being available: a Revelio individual inherits the DIME contribution history of the L2 voter record to which she was matched.

Of the 44,124,662 individuals in the Revelio-L2 matched pool, approximately 14.4% are linked to at least one DIME contributor record. Within the 2012-2022 analytical panel of 380,153 registered partisans employed at PAC-sponsoring firms, 61,300 (16.1%) have a DIME record. This DIME-matched sub-sample is used for: (a) the replication (Section 3.3 Table 4 Columns 1 and 2), which requires a contributor identifier to construct the panel; (b) the validation analysis (Section 3.1), which compares voter-registration partisanship to Li's DIME-based partisanship measure on the overlap sample; and (c) the decomposition analysis (Section 3.2, Table 3), which requires both partisanship measures on the same individuals.

Give = 0 is assigned to all panel observations with no matched DIME donation to the relevant PAC in the relevant cycle, including the 99.12% of panel pair-cycles where no donation occurs. This assignment is the key design innovation relative to Li (2018): it permits the extensive-margin specification by treating non-donation as a genuine observed outcome rather than an absence of data.

A.2 DIME-PAC Crosswalk Coverage

Procedure. Linking eligible employers to the PACs sponsored by their employers requires fuzzy-matching PAC sponsor names from FEC committee records to employer names in the Revelio employment data. I follow the standard string-similarity approach used in the campaign finance literature, using Jaro-Winkler similarity on cleaned PAC and employer names.

Both PAC names and employer names are pre-processed through a common cleaning function: standardizing case, stripping punctuation, removing common corporate suffixes ("Inc," "LLC," "Corporation"), and removing PAC-specific suffixes ("Political Action Committee," "PAC," "Fund"). For each cleaned PAC name, I compute the Jaro-Winkler similarity against every employer name appearing at least 20 times in the L2-linked Revelio data. I retain the closest match per PAC where similarity exceeds a threshold:

1. **Tier 1** matches require similarity ≥ 0.97 (effectively exact-name matches after cleaning)
2. **Tier 2** matches require similarity ≥ 0.88 (close matches accommodating minor variation)
3. Lower-similarity matches are discarded as unreliable

Cleaned PAC names with fewer than four characters after suffix-stripping are excluded to guard against over-stripping (e.g., a PAC name reduced to a generic two-letter abbreviation).

Results. Of Li's (2018) corporate PAC universe of 5,284 access-seeking PACs, 2,162 (40.9%) are matched to an identifiable sponsor firm in the Revelio employment data: 1,161 Tier 1 matches and 1,001 Tier 2 matches. Table A.3 reports the breakdown.

Of the 2,162 matched PACs, 840 appear in the 2012-2022 analytical panel with at least one observed pair-cycle observation. The remaining 1,322 matched PACs are dropped from

Table A.3: DIME–PAC crosswalk coverage

Stage	N	Share
Li (2018) PAC universe[a]	5,284	–
PACs matched to Revelio employer name[b]	2,162	40.9%
Tier 1 (Jaro–Winkler ≥ 0.97)	1,161	53.7%
Tier 2 (Jaro–Winkler $\in [0.88, 0.97)$)	1,001	46.3%
Matched PACs validated (`match_valid=True`)	2,159	99.9%
Matched PACs appearing in 2012–2022 analytical panel	840	38.9%
Matched PACs not in panel[c]	1,322	61.1%

Notes: Coverage of Li's (2018) corporate PAC universe by the DIME–PAC name-matching crosswalk (Jaro–Winkler similarity on cleaned committee and employer names).

[a] Li (2018) reports a corporate PAC universe of 5,284 PACs across the 1990–2016 cycles. The 5,284 denominator is used as reported in the paper and is not independently verifiable from current data. PACs with missing or empty FEC committee names ("513 of 5,284" per the text) cannot be independently recomputed without her universe file.

[b] Total matched counts include three Tier 1 rows with `match_valid=False` that were manually flagged; these are retained in the count of potential matches but excluded from the panel.

[c] PACs sponsored by firms whose employee base does not appear in Revelio during 2012–2022 (dissolved before 2012, inactive in panel window, or too small to register in LinkedIn data).

the analytical sample because they have no employees observable in Revelio during the 2012-2022 window - typically PACs sponsored by firms whose employee base is too small to appear in the Revelio data with any frequency, PACs whose sponsor firms dissolved before 2012, or PACs that have become inactive during the panel window.

A.3 PAC Representativeness

Comparison. The 2,162 crosswalk-matched PACs that feed into the analytical panel constitute 41% of Li's 5,284-PAC corporate universe. Whether the matched subset is representative of the full universe on dimensions relevant to my analysis determines how much inference we can draw about Li's population on this panel. This appendix compares matched to unmatched PACs on three dimensions: partisan lean, fundraising size, and temporal activity. The comparisons are restricted to the 2,947 PACs in Li's universe with any candidate-contribution activity during 2012–2022 (1,224 matched, 1,723 unmatched); PACs from Li's universe that were inactive throughout the panel window cannot contribute to a representativeness test.

Table A.4 reports the matched PACs versus unmatched PACs along key dimensions. The comparison uses Li's partisan lean measure (PAC_R_{jt}) averaged across cycles where each PAC is active, FEC-reported total disbursements across 2012–2022, and the count of cycles during which each PAC reports any candidate contribution activity.

Table A.4: Matched vs. unmatched PAC comparison

	Matched (N=1,224)			Unmatched (N=1,723)			
Dimension	Mean	SD	Median	Mean	SD	Median	p
Partisan lean (PAC_R_{jt})[a]	0.129	0.229	0.121	0.125	0.260	0.130	0.642
Disbursements 2012–2022 ($K)[b]	1316	3140	319	809	2405	155	$<10^{-6}$[c]
Cycles active 2012–2022	4.77	1.69	6.0	4.42	1.81	5.0	$<10^{-7}$

Notes: PAC-level statistics across the 2012–2022 window. Universe is the 2,947 unique committee IDs appearing in `pac_rjt.csv` during 2012–2022 (the set of PACs with computable partisan-lean history in the panel window; this is the closest on-disk analog of Li's 5,284 universe, which is not itself available).

[a] Per-PAC mean across cycles in which the PAC is active.

[b] Per-PAC sum of FEC-reported total disbursements across all six cycles 2012-2022. Source: FEC committee summary bulk files. Disbursements are summed across the six cycles per committee, then restricted to the 2012-2022 universe to keep the scope consistent with the other two dimensions. Sample sizes: $N = 1{,}224$ matched, $N = 1{,}723$ unmatched PACs.

[c] Mann–Whitney U test (robust to the heavy right tail in disbursements): $p < 10^{-19}$. (A mean-based t-test understates the matched–unmatched difference because the disbursement distribution is heavy-tailed; the rank-based test is the recommended inference for this comparison.)

Partisan lean is statistically identical. Matched PACs have a mean PAC_R of 0.13; unmatched PACs have a mean PAC_R of 0.13; the difference is not statistically distinguishable from zero ($p = 0.64$). This is the dimension most relevant to the identification of the alignment effect: if matched PACs were systematically more Republican-leaning than unmatched PACs, the alignment coefficient identified on our panel might not generalize to Li's full universe. The null result on partisan lean rules out this concern.

Matched PACs are larger. Across the full 2012–2022 window, matched PACs have a median total disbursement of \$319K per PAC compared to \$155K for unmatched PACs, and a mean of \$1,316K compared to \$809K (t-test $p = 7.3 \times 10^{-7}$; Mann–Whitney rank-sum test, robust to the heavy right tail of the disbursement distribution, $p < 10^{-19}$). Matched PACs are roughly 2.1× larger by median and 1.6× larger by mean. This is expected: larger sponsor firms have more employees visible in voter registration data and therefore more opportunities for the fuzzy-matching procedure to succeed. The size difference is substantial but one-sided (the panel over-represents larger PACs), not distributional (there are PACs across the full size range represented in the sample).

Matched PACs are more persistently active. Matched PACs are active in a mean of 4.8 cycles across the analytical window, compared to 4.4 cycles for unmatched PACs ($p < 10^{-7}$). This reflects the same underlying pattern as the size difference: larger, more established firms sponsor PACs with longer active histories, and these PACs are more likely to match. The persistence difference is modest (half a cycle on average) but statistically robust.

Implications. The size and persistence biases do not threaten internal validity of the alignment estimate: Spec 3 identifies the alignment coefficient from cross-individual variation within PAC-cycle, which is identified equally well on a large PAC as on a small one. What the biases affect is the external generalizability of the counterfactual dollar magnitude in Appendix A.8, which extrapolates the per-PAC suppressed revenue estimate to Li's full 5,284 PAC universe. Because matched PACs are approximately 1.6× larger than unmatched PACs by mean disbursements (2.1× by median), a per-PAC extrapolation to unmatched PACs likely overstates the total suppressed revenue by 40-60% if suppressed donations scale with PAC size. The counterfactual extrapolation in Appendix A.8 should therefore be read as an upper-bound estimate of the total suppressed revenue in Li's universe.

The partisan-identity null result is the substantively important finding for the identification claim. Because the matched subset of PACs is ideologically representative of Li's universe, the alignment effect estimated on the matched panel is defensibly interpreted as characterizing the partisan response on the broader universe of access-seeking PACs, not just on the selected subset.

A.4 Panel Window Diagnostic

The analytical panel runs from 2012-2022, six biennial federal election cycles. This appendix documents the diagnostic work behind two construction choices that are not obvious from the panel itself: dropping the 2010 cycle from the analysis window despite Revelio employment coverage extending earlier, and dropping the 2024 cycle despite the FEC reporting cycle being complete.

The 2024 cycle: snapshot timing. The 2024 cycle is excluded because the Revelio employment data is observed at an October 2023 snapshot. Cycles ending after October 2023 - including the 2024 cycle, which spans January 2023 through December 2024 - would require imputing employment status for unobserved months. Rather than introduce imputation, I restrict the upper bound of the panel to 2022.

The 2010 cycle: a build-process coverage gap. The 2010 cycle is excluded for a more specific reason. The DIME contribution database used to construct the individual-to-PAC donation records is operationally complete for the 2010 cycle: querying the raw DIME database for itemized contributions ($\geq$\$200) from the target donor population to any PAC committee yields 28,916 donation records in 2010, comparable to 33,803 in 2014 and 47,317 in 2016. The raw donation activity exists.

The problem emerges at the next step in the data pipeline. Of the 28,916 raw 2010 donations, only 7% are retained when I join them to the 2,162 matched PACs. For comparison, the same join in 2012 retains 38% of raw donations; on 2014, 43%; on 2016, 44%. The retention rate for 2010 is more than 5× lower than for any subsequent cycle.

This pattern reflects two interesting features of the underlying data. First, DIME's PAC recipient identifiers are less consistently populated for pre-2012 contributions, reflecting the implementation of the FEC e-filing mandate that took effect in 2012 and standardized committee identifier reporting from that point forward. Many 2010 donations exist in DIME but cannot be matched to the PAC universe because their committee identifiers do not align with the post-2012 standard.

Second, the target donor cohort is heavily skewed post-2012 by construction. Of the 224,679 individuals matched between Revelio employment records, L2 voter registration, and DIME contribution records, only 10.3% have their first observed DIME contribution on or before 2010. The remaining 89.7% are either first active in 2012 or later, or are first active in cycles before 2010 but with their contribution histories concentrated post-2010. This cohort tilt is itself a function of the underlying data: the Revelio employment records that anchor our panel construction are sparse before approximately 2010, and the L2 voter registration linkage is more reliable for individuals whose employment spells fall in the late-2000s and after.

The combination produces a 2010 cycle with implausibly low donation event counts in the analytical panel: 265 pair-cycle observations with Give=1, compared to 1,628 in 2012 and 1,247 in 2022. Even with a corrected build pipeline that recovered the lost 93% of 2010 donations, the cohort tilt would leave the 2010 cycle with substantially fewer donor-cycle observations than later cycles. The cycle is unrecoverable to a level of coverage consistent with 2012-2022 without rebuilding the underlying employment-voter-donor linkage upstream.

Verifying that 2022 is clean. Before committing to the 2012-2022 window, I verified that the 2022 cycle does not exhibit a coverage problem analogous to 2010. The diagnostic compares raw DIME donation activity to the analytical file across cycles. The 2022 cycle has 62,908 raw donation records from target donors (compared to 47,317 in 2016 and 45,910 in 2018) and 1,247 records retained in the analytical file after the PAC-universe join. The retention rate (1,247/62,908 = 2.0%) is lower than 2014-2018 (3.4-4.8%) but the absolute donation count is comparable - the lower retention reflects a denominator effect from the 2020-2022 small-donor wave that inflated raw donation counts without proportionally increasing donations to the access-seeking PACs in this universe. This is real campaign-finance behavior, not a data-quality issue, and the 1,247 retained donations support cleanly identified estimates.

Effect on headline coefficients. Restricting from a 2010-2022 panel to a 2012-2022 panel reduces the analytical sample from 239,467 to 205,522 pair-cycle observations and increases the pooled extensive-margin coefficient from $\hat{\delta} = 0.0303$ ($p = 1.7 \times 10^{-8}$) to $\hat{\delta} = 0.0352$ ($p = 1.3 \times 10^{-8}$). The 16% increase in the pooled coefficient reflects the removal of the 2010 cycle, which contributed near-zero per-cycle estimates - a mechanical consequence of the sparse 2010 coverage that pulled the pooled estimate downward without corresponding to a substantively different mechanism. The substantive finding is robust to either window choice: the extensive-margin coefficient is highly significant and substantively large on both the 2010-2022 panel ($\hat{\delta} = 0.0303$, $p = 1.7 \times 10^{-8}$) and the 2012-2022 panel ($\hat{\delta} = 0.0352$, $p = 1.3 \times 10^{-8}$). The 2012-2022 window provides a cleaner estimate of the aligned effect than the 2010-2022 window because it excludes the documented 2010 build-process coverage gap, and is the basis for the headline result reported in Section 3.3.

A.5 Modified Granger Causality Test

Specification. Following Li (2018, Appendix A.2.5), I assess the parallel-trends assumption underlying Equation 1 via a modified Granger causality test. The specification augments the baseline intensive-margin estimator with two lags and three leads of the alignment interaction:

$$
\begin{aligned}
\text{Give}_{ijt} = \alpha_{ij} + \tau_{jt} \\
&+ \beta \left(\text{Indiv\_R}_i^{VR} \times \text{PAC\_R}_{jt}\right) \\
&+ \sum_{m=1}^{2} \gamma_m \left(\text{Indiv\_R}_i^{VR} \times \text{PAC\_R}_{j,t-m}\right) \\
&+ \sum_{q=1}^{3} \lambda_q \left(\text{Indiv\_R}_i^{VR} \times \text{PAC\_R}_{j,t+q}\right) + e_{ijt}
\end{aligned}
\tag{3}
$$

retaining the donor-PAC fixed effect α_{ij} and PAC-cycle fixed effect τ_{jt}, with standard errors clustered at the PAC level. The parallel trends assumption requires the joint null $H_0 = \lambda_1 = \lambda_2 = \lambda_3 = 0$. Under the null, observed donation patterns change only when PAC partisan lean actually changes, not in anticipation of future changes.

Sample restriction. The lead structure imposes a mechanical sample restriction: estimating Equation 3 requires PAC_R values at cycles $t + 1$, $t + 2$, and $t + 3$, which constrains the effective sample to cycles for which three subsequent observations exist in the panel. With the analytical window ending at 2022, only cycles 2012, 2014, 2016, and 2018 contribute identification - cycles 2020 and 2022 drop because their lead-2 and lead-3 values would require PAC_R for 2026 and 2028. This is a mechanical endpoint loss inherent to leads-and-lags estimation, not a substantive restriction. I report this transparently because it has consequences for interpreting the current-period coefficient (see discussion below).

Results. I estimate Equation 3 on two samples: the DIME-matched pure-partisan subsample used in the main analyses (Column 1) and the full analytical panel (Column 2).

Table A.5 reports the coefficients and standard errors.

Table A.5: Modified Granger causality test: Lead and lag coefficients

	(1) DIME-matched	(2) Full panel
Lag 2 ($t-2$)	+0.00492	+0.00035
	(0.00680)	(0.00105)
Lag 1 ($t-1$)	−0.00800	−0.00055
	(0.00832)	(0.00107)
Current (t)	+0.01764†	+0.00245†
	(0.00916)	(0.00125)
Lead 1 ($t+1$)	+0.00723	+0.00119
	(0.00993)	(0.00127)
Lead 2 ($t+2$)	−0.00235	−0.00070
	(0.00912)	(0.00145)
Lead 3 ($t+3$)	+0.00174	+0.00056
	(0.00677)	(0.00114)
Wald $H_0 : \lambda_1 = \lambda_2 = \lambda_3 = 0$	$\chi^2(3) = 1.356$, $p = 0.716$	$\chi^2(3) = 2.593$, $p = 0.459$
N obs	102,358	498,004
N unique pairs	45,603	267,027
N PACs	593	715

Notes. Li (2018) A.5 modified Granger test. Pair + PAC×cycle FE; PAC-clustered SEs. Cycles 2012-2022 (six cycles); cycles 2020 and 2022 drop out of the regression because lead-2 and lead-3 require PAC_R values beyond 2024, which are unavailable in 'data/pac_rjt.csv'. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

The joint Wald test fails to reject the null of parallel trends in both samples ($\chi^2(3) = 1.36$, $p = 0.72$ on the DIME-matched sample; $\chi^2(3) = 2.59$, $p = 0.46$ on the full panel). None of the three lead coefficients is individually significant at conventional levels, and their confidence intervals straddle zero. By comparison, Li (2018) reports analogous Wald test p-values of 0.139 (pure partisans) and 0.218 (all partisans) on her panel; the larger p-values here are more clearly consistent with the parallel trends assumption.

Figure A.1 plots the lead and lag coefficients with 95% confidence intervals across both samples. Pre-treatment leads are tightly centered on zero in both columns, providing visual confirmation that observed donation patterns do not respond in anticipation of future PAC partisan shifts.

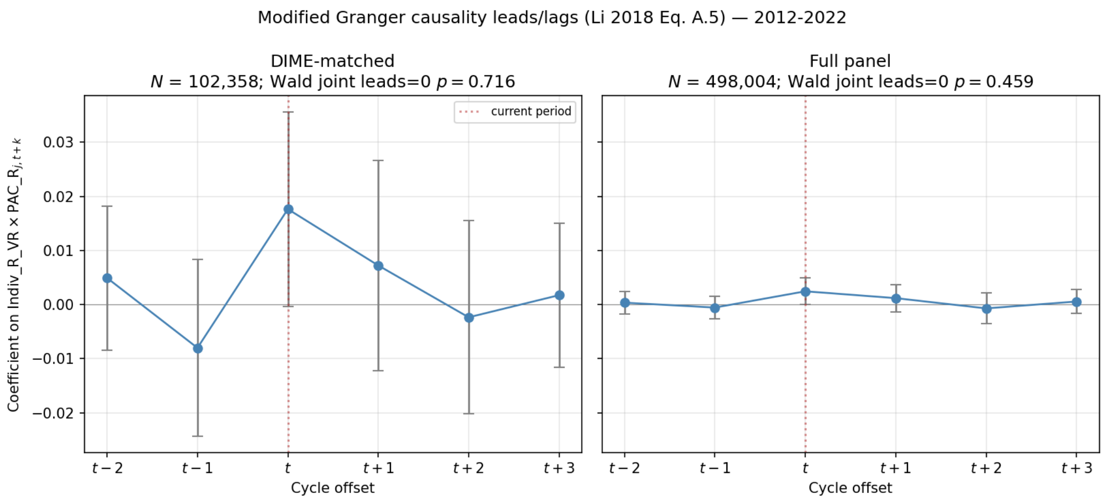

Figure A.1: Modified Granger causality test: Lead and lag coefficients with 95% CIs
Note: Coefficients on the alignment interaction (Indiv_$\mathrm{R}_i^{VR} \times$ PAC_R_{jt}) at $t-2$ through $t+3$ from Equation 3, with 95% confidence intervals. Dashed horizontal line at zero. Filled circle marks the current period (t). Standard errors clustered at PAC level. Left panel: DIME-matched subsample. Right panel: full analytical panel.

Interpretation: current-period coefficient. The current-period alignment coefficient on the DIME-matched sample in Equation 3 is $\hat{\beta} = +0.01764$ (SE = 0.00916), larger than the baseline Spec 1 estimate of $\hat{\beta} = +0.00589$ on the full 2012-2022 window. This difference reflects the sample restriction documented above: the Granger regression is identified off the four cycles 2012-2018, while baseline Spec 1 uses all six cycles 2012-2022. Restricting to 2012-2018 captures the upward portion of the per-cycle alignment trajectory ($\hat{\delta}_{2012} = 0.030$ rising to $\hat{\delta}_{2018} = 0.045$, per Figure 1) without the 2022 attenuation ($\hat{\delta}_{2022} = 0.024$) that pulls the pooled estimate down. The discrepancy is a mechanical consequence of the panel cycle subset, not evidence of internal inconsistency between the Granger specification and the baseline. The full-panel current-period coefficient is $\hat{\beta} = +0.00245$, of similar order of magnitude to the full-panel estimate from related specifications.

The modified Granger test finds no evidence against parallel-trends assumption on either sample. Lead coefficients are jointly and individually consistent with zero; the joint Wald p-values exceed Li's reported benchmarks. The baseline estimates in Section 3 are not confounded by anticipatory dynamics in donor behavior with respect to future PAC partisan shifts.

A.6 Calibrated Monte Carlo Simulation

Section 3.2 establishes that Li's within-donor coefficient of $\hat{\beta} = 0.0853$ is recovered on the ever-donor sub-sample but not on the broader DIME-matched sample. A natural concern is that this recovery is mechanical: restricting to pairs with at least one observed donation selects pairs with multiple Give events, providing within-pair variation that the never-donor majority does not provide. Under this view, Li's estimated magnitude could reflect a small underlying behavioral parameter inflated by selection on the ever-donor outcome.

This appendix documents the calibrated Monte Carlo simulation that rules out this explanation. The simulation asks: if the true within-donor parameter were small (set to 0.02), could selection on the ever-donor outcome mechanically generate coefficients of Li's magnitude when Equation 1 is estimated on the simulated ever-donor sub-sample? The answer is no.

I simulate pair-cycle panels matching the cycle count, PAC count, and alignment distribution of the analytical sample:

1. Six cycles, mirroring the 2012-2022 panel window
2. 2,000 PACs, with PAC partisan lean PAC_R_{jt} drawn from the empirical distribution of PACs for the 2012-2022 cycles
3. 50 individual-PAC pairs per PAC (100,000 total pairs; 600,000 pair-cycles per draw)
4. Individual partisanship $\text{Indiv\_R}_i \in -0.5, +0.5$ assigned with equal probability
5. Donation outcomes generated by a fixed-effects logit data-generating process (DGP) with within-pair partisan response parameter set to 0.02 — substantially smaller than Li's reported 0.0853. (Note: the simulation parameter governs the within-pair effect identified by Equation 1, corresponding to β in the main-text notation, not the cross-individual δ from Equation 2. We use a generic symbol below to avoid notation collision with the main-text δ.)

6. Pair fixed effects drawn from an empirical distribution that produces a baseline participation rate p_0, varied across the simulation sweep

For each value of p_0, I draw 100 simulated panels and estimate Equation 1 (Li's specification with pair and PAC-cycle fixed effects) on two samples: the full simulated panel and the simulated ever-donor sub-sample (pairs with at least one Give = 1 across the simulated cycles). The full-panel estimates serve as a sanity check that the DGP recovers the true parameter ($\delta = 0.02$) when no sample restriction is applied; the ever-donor estimates are the quantity of interest.

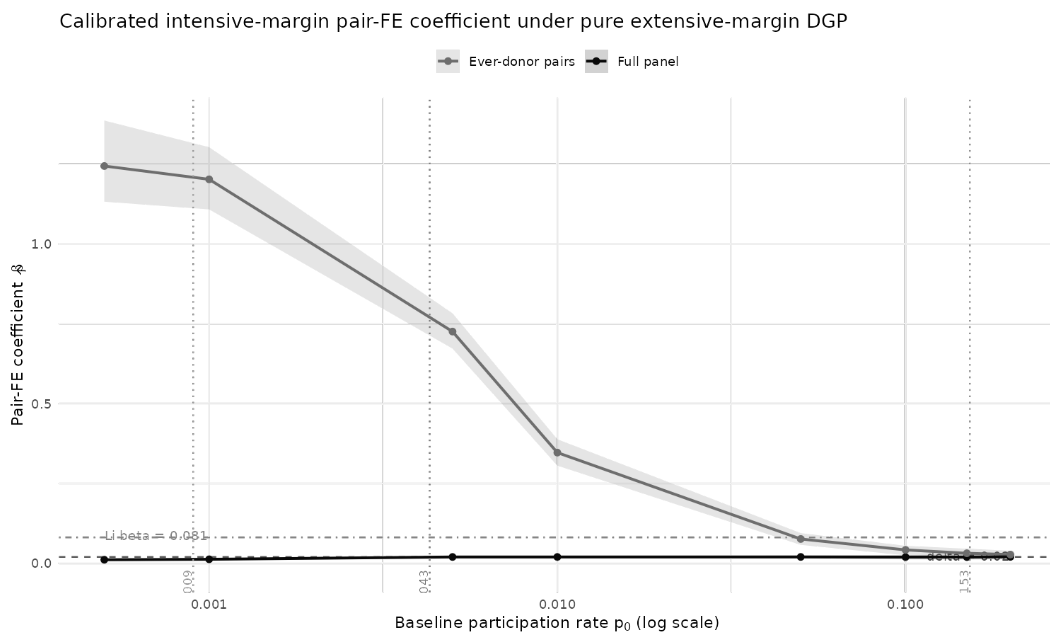

Figure A.2: Calibrated Monte Carlo simulation: Full-panel and ever-donor coefficients vs. participation rate

Note: Simulated coefficient estimates from Equation 1 across values of baseline participation rate p_0. Full-panel estimates shown in one series; ever-donor estimates in another. Li's (2018) reported coefficient of $\hat{\beta} = 0.0853$ shown as horizontal reference line. 100 draws per p_0 value.

Results. Figure A.2 plots the simulated coefficient estimates across the range of p_0 values. Two patterns are evident. First, the full-panel estimate stabilizes at approximately $\hat{\beta} \approx 0.020$ across realistic values of p_0, confirming that the DGP recovers the true parameter (0.02) when no sample restriction is applied. Second, the ever-donor estimate exhibits the signature pattern of selection-on-outcome bias: at very low participation rates ($p_0 < 0.001$), the ever-donor coefficient is highly inflated (above 1.0) because the selected sub-sample is dominated by a handful of high-propensity pairs; as p_0 increases, the ever-donor coefficient declines monotonically toward the true parameter.

At $p_0 = 0.15$ - Li's reported participation rate in her sample - the simulated ever-donor coefficient has a mean of 0.0316 and a standard deviation of 0.0067 across the 100 draws. Li's reported $\hat{\beta} = 0.0853$ falls 8.0 standard deviations above this mean, and not a single draw across the 100 produces a coefficient as large as 0.0853. Across the entire range of p_0 values examined, the simulated ever-donor distribution does not reach Li's coefficient at any participation rate consistent with her sample.

Interpretation. Selection on outcome alone, applied to a small underlying behavioral parameter, cannot mechanically generate a coefficient of Li's magnitude. The within-donor partisan response on the ever-donor sub-population that Li identifies is genuinely large; it is not an artifact of restricting the sample to observed donors.

This result complements rather than substitutes for the empirical decomposition in Section 3.2. The decomposition shows that on real data, the entire gap between the broader sample coefficient ($\hat{\beta} \approx 0.006$) and the ever-donor coefficient ($\hat{\beta} = 0.21$) is driven by the ever-donor sample restriction. The simulation shows that this gap is too large to be explained by mechanical selection inflation under a small true parameter. Together, they support the substantive interpretation: ever-donors are behaviorally distinctive in their within-pair responsiveness to alignment shifts, and the magnitude Li identifies reflects this real difference rather than a methodological artifact of her sample construction.

A.7 Placebo Test

Specification. To assess whether the extensive-margin alignment effect identified in Equation 2 could be a chance artifact of the regressor structure, I conduct a placebo test based on within-PAC permutation of voter registration. For each of 500 placebo draws, I randomly shuffle Indiv_R_i^{VR} across individuals within each PAC, breaking the alignment relationship between employee partisanship and PAC partisan lean while preserving the marginal distribution of partisanship within each PAC. I then re-estimate Equation 2 on the shuffled data and store the resulting alignment coefficient.

If the observed coefficient $\hat{\delta} = 0.0352$ reflects a genuine partisan response, the placebo distribution should be centered at zero with a small standard deviation, and the observed coefficient should fall far above the placebo mass. If the observed coefficient instead reflects an artifact of the fixed-effects structure or distributional properties of the data, the placebo distribution should overlap with 0.0352.

Results. The placebo distribution is centered at $\bar{\delta}_{\text{placebo}} = 0.000003$ with a standard deviation of 0.0026. The observed coefficient of 0.0352 falls 13.4 standard deviations above the placebo mean. Across the 500 placebo draws, the maximum coefficient observed is well below the true estimate; not a single placebo draw produces a coefficient as large as 0.0352 (empirical $p < 0.002$).

Figure A.3 plots the placebo distribution alongside the observed coefficient. The placebo mass lies entirely within ± 0.005, while the observed estimate sits at $+0.0352$, well outside the placebo range.

The result rules out the possibility that the Spec 3 alignment coefficient reflects mechanical features of the panel structure - clustering, fixed-effects absorption, or the marginal distribution of partisanship - that would produce non-zero coefficients under any assignment of Indiv_R_i^{VR} to individuals within PACs. Whatever the Column 3 coefficient is measuring, it is identified from the actual partisan-alignment relationship between employees and PACs, not from noise in the alignment construction.

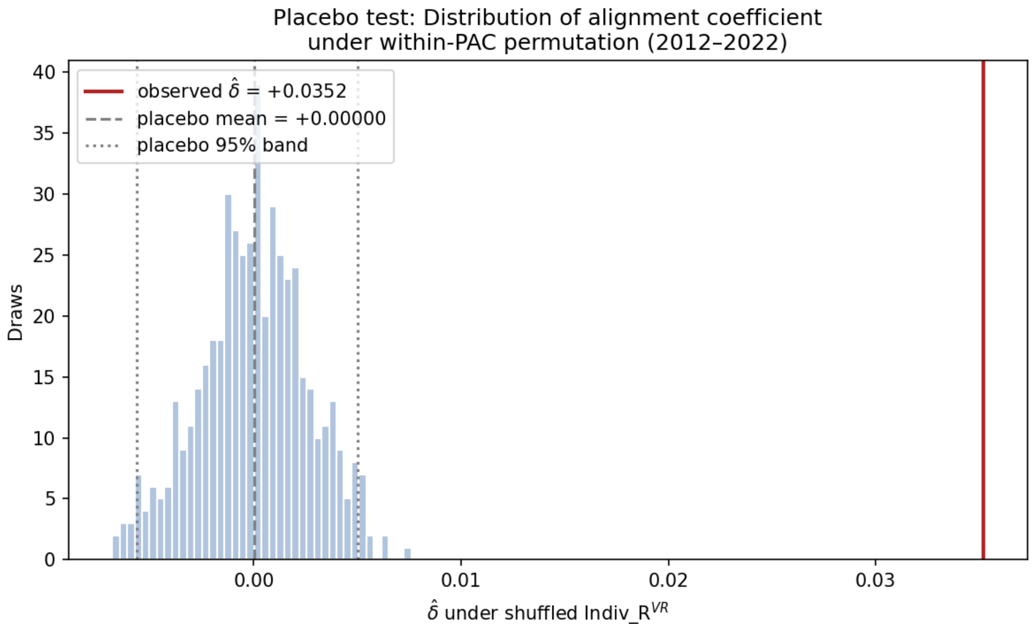

Figure A.3: Placebo test: Distribution of alignment coefficient under within-PAC permutation

Note: Histogram of alignment coefficients from 500 placebo draws in which voter registration is randomly shuffled within each PAC, breaking the alignment relationship. Observed coefficient of $\hat{\delta} = 0.0352$ marked. Placebo mean: 0.000003; standard deviation: 0.0026.

A.8 Counterfactual Revenue Suppression

The extensive-margin coefficient $\hat{\delta} = 0.0352$ identified in Section 3.3 quantifies the per-cycle probability of donation that a one-unit shift in alignment produces. Translating this probability shift into a dollar magnitude requires computing the donation revenue suppressed across the panel by partisan misalignment between eligible employees and the PACs sponsored by their employers. This appendix documents the calculation and the assumptions on which it relies.

Calculation. For each pair-cycle observation (i, j, t) in which alignment is negative - that is, PAC_R$_{jt}\times$ Indiv_R$_i^{VR} < 0$ - partisan misalignment suppresses the per-cycle probability of donation by $\hat{\delta}\times|\text{Alignment}_{ijt}|$. Aggregated across misaligned pair-cycles, the total suppressed donation events per cycle on the matched panel is:

$$\widehat{N}^{\text{suppressed}} = \frac{1}{N_{\text{cycles}}} \sum_{(i,j,t):Alignment_{ijt}<0} \hat{\delta} \times |Alignment_{ijt}| \tag{4}$$

where $N_{cycles} = 6$ for the 2012-2022 window. Multiplying by the mean donation amount conditional on giving converts events to dollars:

$$\widehat{D}^{\text{suppressed}} = \widehat{N}^{\text{suppressed}} \times \bar{d} \tag{5}$$

where $\bar{d}$ is the mean donation amount among observed Give = 1 events on the panel.

Inputs. The relevant inputs from the 2012-2022 panel are:

1. $\hat{\delta} = 0.0352$ (SE = 0.0062), the extensive-margin coefficient from Equation 2
2. 112,950 misaligned pair-cycle observations (54.96% of the 205,522-observation panel)
3. 89,476 aligned pair-cycle observations (43.54%); 3,096 zero-alignment observations (1.51%, occurring when PAC_R$_{jt} = 0$)

4. Mean $|\text{Alignment}_{ijt}| = 0.0644$ among misaligned pair-cycles; sum of $|\text{Alignment}_{ijt}|$ across misaligned pair-cycles = 7,271.73

5. Mean donation amount conditional on giving: $\bar{d} = \$2,525$ (median \$1,000, SD = \$9,191 across 9,281 observed Give = 1 events with positive amounts)

The donation amount distribution exhibits a heavy right tail; I report sensitivity to using the median in place of the mean below.

Matched-panel suppressed revenue. Applying Equation 4 to the 2012-2022 panel yields 42.7 suppressed donation events per cycle on the matched panel of 840 PACs. Multiplying by $\bar{d} = \$2,525$ yields \$107,703 in suppressed donation revenue per cycle. The 95% confidence interval, propagating the standard error by $\hat{\delta}$ via the delta method, is \$70,522 to \$144,884 per cycle.

Universe extrapolation. The matched-panel estimate covers 840 of Li's 5,284 access-seeking PACs (15.9%). To produce a universe-scale estimate, I scale the per-PAC suppressed revenue from the matched panel (\$128 per PAC per cycle) to the full 5,284 PAC universe, yielding \$677,502 in suppressed donation revenue per cycle for Li's full universe (95% CI: \$443,614 to \$911,390).

The extrapolation rests on the assumption that matched and unmatched PACs face similar misalignment distributions among their employee populations. Appendix A.3 establishes that matched and unmatched PACs are statistically indistinguishable on partisan lean ($p = 0.64$), supporting this assumption on the dimension most relevant to the calculation. The matched-vs-unmatched difference on PAC size (matched PACs approximately 1.6× larger by mean disbursements, 2.1× by median) suggests the unmatched PACs have smaller employee populations, in which case the per-PAC scaling overstates the unmatched contribution. The universe should therefore be read as an upper bound on the suppressed revenue Li's mechanism implies at the population scale.

Sensitivity to donation amount. The mean donation amount of \$2,525 is heavily influenced by the right tail of the donation distribution; the median is \$1,000. Using the median in place of the mean scales the dollar suppression estimates by a factor of 0.396:

1. Matched-panel suppressed revenue per cycle: \$42,651 (vs. \$107,703 using the mean)

2. Universe-extrapolated suppressed revenue per cycle: \$268,357 (vs. \$677,502 using the mean)

Both the mean-based and median-based estimates support the substantive conclusion that the partisan-alignment mechanism shapes a population-scale fundraising pattern of meaningful magnitude. The mean is the appropriate central estimate for the dollar-suppressed quantity (which is the sum of dollars not given, not the typical pair-cycle's contribution); the median provides a more conservative lower bound that excludes the influence of the largest donations.

Sanity check. The calculation predicts that observed donation events should be more concentrated among aligned pair-cycles than under a null in which alignment plays no role in participation. With aligned pair-cycles representing 43.5% of the panel and misaligned pair-cycles representing 55.0%, the null prediction (no alignment effect) would produce a 43.5% / 55.0% / 1.5% split among observed Give events. The actual split is 54.8% aligned (989 events), 44.7% misaligned (807 events), 0.6% zero (10 events) - substantively over-representing aligned pair-cycles relative to their share of the panel. This is the qualitative signature of the extensive-margin mechanism the Spec 3 coefficient is identifying.

Table A.6: Counterfactual revenue suppression: Matched panel and universe extrapolation

Sample	Donor-events / cycle	95% CI	Dollars / cycle	95% CI
Matched panel (840 PACs)	42.7	[27.9, 57.4]	\$107,703	[\$70,522, \$144,884]
Per matched PAC	0.1	–	\$128	–
Universe-extrapolated (5,284 PACs)	268.4	[175.7, 361.0]	\$677,502	[\$443,614, \$911,390]

Inputs: $\hat{\delta} = 0.0352$ (SE 0.0062); $\bar{d} = \$2,524.63$; $\sum_{\text{misaligned}} |\text{Alignment}| = 7,271.7$; $N_{\text{cycles}} = 6$.