# Money Follows the Winner: Office Premiums and Access-Seeking Across the Donor Spectrum in U.S. House Elections

Kisoo Kim*

August 25, 2026

**Abstract**

Does campaign giving buy access, or express policy preferences? I apply a regression discontinuity to the win-loss margin from the donor's perspective using 1.8 million committee-by-race and 20.4 million individual-by-race observations for U.S. House elections, 1990–2024. A committee is 20 percentage points likelier to keep funding its pick next cycle when the pick holds the seat, and 7 points likelier to fund the opponent when the opponent does: office itself carries a premium on both relationships. Loyalty to defeated candidates falls, total giving is unchanged, and portfolios tilt toward whichever party wins, the joint signature of access-seeking. Among publicly traded firms' committees, lobbying directed at the winner, but not at the loser, shows the same discontinuous increase. Ideological committees follow too; individual donors follow in proportion to the scale of their giving, small donors barely at all. Organized and large-scale giving buys access; small donors express preferences.

**Keywords:** campaign finance; interest groups; access; regression discontinuity; U.S. Congress; lobbying

*Email: kisoo.kim@vanderbilt.edu

# 1 Introduction

In the 2014 election cycle, 592 committees gave $3.0 million to the reelection campaign of Pete Gallego, a first-term Democrat from Texas's Twenty-Third District. That November, Gallego lost the seat to Republican Will Hurd, 47.7 percent of the vote to 49.8. The next cycle was a rematch: Gallego ran to take the seat back in 2016, now as a challenger against the incumbent. The committees that had financed Gallego gave Hurd $912,000 and Gallego's comeback $733,000, and among the business committees in that group the switch was nearly complete: of the fourteen that funded either candidate again, thirteen funded Hurd, $89,000 against $1,000. The pattern recurs cycle after cycle: the committees that finance U.S. House campaigns behave as if the office matters more than its occupant. What that behavior reveals about the motive behind organized giving has framed three decades of debate.

One view holds that contributions purchase legislative access (face time, information, and attention from whoever holds the seat), a tradition running from Welch (1980) and Cassie and Thompson (1998) to Phillips (2024). A second view, articulated in Ansolabehere et al. (2003) and reflected throughout the cross-sectional CFscore literature (Bonica, 2014, 2018), treats contributions as expressions of ideology: donors back candidates whose positions match their own. A third combines both: committees expand their budgets to cultivate the new winner while maintaining old commitments. The three views collide the moment an election flips the office-holding identity, and the literature has lacked a design clean enough to put them side by side.

When the candidate a committee supported (its pick, in what follows) narrowly loses a U.S. House election, what does the committee do next? The question admits a sharp answer because the three views imply distinct sign patterns across four behavioral outcomes, measured in the election cycle after the race in question (the focal race): whether the committee funds the focal-race winner, whether it continues funding the loser, and, across its entire portfolio of House giving, whether the total amount changes

and whether the partisan composition shifts toward the winner's party. Access-seeking predicts within-budget substitution toward the winner, ideology predicts behavioral continuity, and the hybrid predicts redirection financed by budget expansion. Only the joint signature across all four separates the three.

I estimate the effect of the focal-race outcome on the committee's subsequent giving with a sharp regression discontinuity in the win-loss margin, signed from the donor's perspective: the margin is positive when the committee's pick wins and negative when the pick loses. As-if-random assignment of close-race outcomes (Lee, 2008; Calonico et al., 2014) generates quasi-experimental variation in whether the committee's pick holds office at all. The data cover both contributor universes of the Database on Ideology, Money in Politics, and Elections (Bonica, 2014): 1.8 million committee-by-race observations (cycles 1990 through 2022, outcomes measured in the following cycle) and 20.4 million individual-by-race observations built under identical race definitions and pick rules. The discontinuity is a between-committee contrast, and a within-committee event study confirms the redirection directly.

I find that the outcome of the election reshapes committee contributions in the next election cycle. Office carries a premium on both of a committee's relationships in the race. Among committees whose pick narrowly lost, one in four still fund that pick; among those whose pick narrowly won, the share is 20 percentage points higher (the continuation premium). Among committees whose opponent narrowly lost, fewer than one in twenty fund that opponent next time; among those whose opponent narrowly won, the share is 7 points higher (the crossover premium): money that opposed a candidate turning to support them because they won. Both premiums are paid where the defeated candidate leaves the seat's ballot, as two-thirds of narrow losers do; where the loser runs again, backers on both sides stay with their candidate at the same rate and neither side crosses over. Total giving is unchanged, and portfolios' partisan composition shifts toward whichever party wins by 6.2 percentage points (Section 4.1).

The pattern matches the access-seeking signature and only that signature. The crossover premium is the component ideology cannot generate: the opponent's proximity to the committee is the same on both sides of the cutoff before the vote, and narrow winners do not moderate in office (Lee et al., 2004; see also Appendix B.5), so money can follow the win only if office itself carries value. The null on total giving falsifies the budget-expansion hybrid: on the margins that expansion would move, totals change by roughly one percent, so the redirection is financed within a near-fixed budget. The winnerward partisan shift moves toward whichever party wins, the opposite of what partisan loyalty predicts. And access does not operate alone: the loyalty and substitution channels identify a proximity weight, ideological or relational, that the model in Section 2 includes explicitly.

A within-design check addresses the default-to-winner reading, on which the committee redirects to the focal winner only because no ideologically similar alternative exists. Replacements are widely available and do receive redirected giving, but giving redirected to the focal winner is more than twice the giving redirected to those replacements (a secondary channel the design cannot decompose into ideology or network ties). A geographic check rules out the state-level counterpart: if what the committee valued were access to the state's House delegation rather than to the focal seat itself, a committee whose pick lost would raise its giving to the winners of the same state's other races; giving to those non-focal in-state winners shows no discontinuity. Committees follow the person who took the seat, not the state.

The paper offers four contributions. First, I provide the cleanest causal evidence to date that committee giving is access-driven, identified at a margin upstream of the committee-assignment thresholds studied in earlier power-following work (Powell and Grimmer, 2016; Fouirnaies and Hall, 2018). I also decompose what prior close-election estimates conflate: continuation with one's own successful pick and genuine crossover to a former opponent. The closest comparison, Fowler et al. (2020), runs the same design

on the same style of panel but measures firm value rather than subsequent giving; their tightly bounded null and my behavioral redirection are jointly consistent with committees purchasing legislative attention, not cash-back policy favors. And on the subsample of committees sponsored by publicly traded firms, where lobbying can be matched to legislators, the redirection is accompanied by a parallel, winner-specific jump in directed lobbying, the activity through which access is exercised, so on that subsample a second behavioral channel supports the access reading.

Second, the four-outcome joint signature is novel in the close-election literature on contributions, whose donor-side studies track individual donors without distinguishing their motivations (Dumas and Shohfi, 2018), study non-partisan municipal contexts (Rueda and Ruiz, 2025), or measure firm value rather than behavior (Akey, 2015; Fowler et al., 2020); with a single outcome, any of the three views can rationalize an estimated discontinuity. The joint design converts two would-be non-results into evidence: the total-giving null establishes within-budget substitution, and the winner-aligned sign convention on the pre-specified composition outcome turns a within-pool cancellation artifact into a direct pooled falsification of partisan loyalty.

Third, I adjudicate access against ideology with two heterogeneity tests that trade identification against substantive strength. The race-side moderator, the winner-loser CFscore gap fixed before the close outcome materializes, shows redirection essentially flat across the race's own polarization; this is the cleanest test the Marshall (2024) critique allows, though of the weaker ideology prediction. The PAC-side moderator tests the stronger prediction, the committee's own ideological cost of funding a distant winner, and shows a tercile-match pattern consistent with a layered ideology channel, but it conditions on a winner trait and is not point-identified under Marshall's bounds. Read jointly, the two tests support an access-dominant reading that depends on neither alone.

Fourth, I show that winner-following is not a property of corporate status: it runs across every organized committee class and rises with individual donation scale, relocat-

ing the access-versus-expression debate from donor type to the scale and organization of giving. The committees an expressive account would call least access-driven (ideological and single-issue committees under a direct purpose classification) redirect at least as strongly as business committees and share the conversion margin. Individual donors locate the divide the committee split cannot: winner-following rises monotonically in donation size, from +0.14 under $200 to +0.47 over $500; partisan reorientation toward the winning party rises from roughly zero to two-thirds of the committee level; and, unlike committee budgets, which are reallocated rather than grown, individual giving responds to the outcome itself, since victory retains and enlarges it. What separates even the largest individual donors from organizations is conversion: money that opposed a candidate turning to support them is +0.069 among committees, +0.011 among large donors, and zero among small ones.

The redirection also changes over time. It is largest before the Bipartisan Campaign Reform Act (+0.46), attenuates after it, and remains substantial through the most recent cycles (+0.32 to +0.38); the attenuation parallels the roughly 50 percent post-2000 decline in donor punishment of extremism that Myers (2026) documents.

## 2 Theory

The committees that fund U.S. House campaigns (political action committees alongside party and candidate committees) pool contributions and disburse them to federal candidates. Their behavior has been read in three competing ways, each with sharp implications for what should happen when a supported candidate narrowly loses an election.

The access view, rooted in a literature that distinguishes access-oriented donor strategies from electoral-ideological ones (Welch, 1980; Cassie and Thompson, 1998) and developed through three decades of empirical work on PAC giving (Romer and Snyder, 1994; Brunell, 2005; Fouirnaies and Hall, 2014, 2018; Barber, 2016; Barber and Eatough, 2020;

Phillips, 2024), treats contributions as purchases of a share of a resource-constrained office's attention: access-oriented donors work with whomever holds the seat and cultivate relationships with both parties (Phillips, 2024; Brunell, 2005). The access value attaches to the office rather than to the legislator personally. On the lobbying side, Hall and Deardorff (2006) model lobbying as a subsidy flowing to *already-aligned* legislators; the access view I test differs in admitting cross-aisle transactions when office changes hands (Brunell, 2005).

The marginal access return to a defeated candidate is small (she neither votes nor sits on committees), and the return to the new winner is large for the symmetric reason. Aldrich et al. (2017) show that leadership-PAC giving responds to majority-party shifts in this way, and Powell and Grimmer (2016) and Fouirnaies and Hall (2018) document analogous power-following at the committee margin, but not at the prior margin my design isolates, the office-versus-no-office threshold itself.

The ideology view, articulated in Ansolabehere et al. (2003) and reflected throughout the cross-sectional CFscore literature (Bonica, 2014, 2018), treats contributions as expressions of policy preferences: aggregate giving is too small to be rational policy investment, so contributions function like consumption, support for candidates whose positions match the donor's rather than investments seeking a return.

A pure ideology account predicts continuity at the win-loss threshold: preferences do not change when an election outcome changes, so a PAC whose pick barely loses should continue giving to that candidate or to ideologically similar co-partisans, not redirect to the new winner. Meisels (2026) shows corporate donors punish extremist nominees more than individual donors do, prioritizing viable power-holders over self-expression; my design extends the contrast to the post-election margin.

A hybrid account combines both motives: ideological commitments keep giving to the loser from collapsing, while access-seeking adds new giving to the winner, which requires expanding the overall budget because the two legs operate on different recipients. Not

formalized in the literature, the hybrid is the natural fallback for any reading in which both motives coexist (Liu, 2022; Stratmann, 2002, 2005).

The views translate into sign predictions on four co-primary outcomes measured at cycle t+2, the next House cycle: giving to the focal-race winner, continued giving to the loser, the change in log total House giving, and the *winnerward shift* in the partisan composition of giving, the sign-aligned form of the pre-specified Democratic-share change (Section 3). No two views share the same complete sign pattern; Table 1 collects the predictions.

The three views can be written as one constrained-allocation problem whose parameters map directly onto the four outcomes, making precise both what the joint signature identifies and what it leaves open. A committee disburses a fixed budget $B$, set in advance by member dues and organizational constraints, across House candidates $c$, choosing nonnegative contributions $g_c$ subject to $\sum_c g_c \leq B$. Each candidate yields a return that combines an access component and an ideological component, scaled by a strictly concave function $\psi$ of the amount given ($\psi' > 0$, $\psi'' < 0$):

$$\max_{\{g_c \geq 0\}} \sum_c \left(V_O\, o_c + V_I\, \rho_c\right) \psi(g_c) \quad \text{subject to} \quad \sum_c g_c \leq B.$$

Here $o_c \in \{0, 1\}$ indicates that candidate $c$ holds office, $\rho_c \in [0, 1]$ is the ideological proximity of $c$ to the committee, and $V_O, V_I \geq 0$ are the committee's valuations of access and of ideological expression. The three views are nested special cases: pure access sets $V_I = 0$, pure ideology sets $V_O = 0$, and the budget-expansion hybrid relaxes the constraint so that $B$ itself rises when the preferred candidate loses.

Let $\omega_c \equiv V_O\, o_c + V_I\, \rho_c$ denote candidate $c$'s allocation weight. The first-order condition equates the marginal return $\omega_c\, \psi'(g_c)$ across funded candidates to the budget's shadow price: the committee funds $c$ if and only if $\omega_c$ clears a threshold, and funds it more generously the larger the weight. The budget thus governs the *scale* of giving while the weights govern its *composition*; an election that leaves $B$ unchanged can still reshuffle the entire

**Table 1:** Four-outcome predictions by model.

| Model | Gives to winner | Gives to loser | Total giving | Winnerward shift |
|---|---|---|---|---|
| Pure ideology | 0 | 0 | 0 | 0 |
| Pure access | + | - | 0 | + |
| Budget-expansion hybrid | + | 0 | + | + |
| Budget-constrained access-dominant | + | - | 0 | + |
| Estimated discontinuity | +0.375 | -0.227 | +0.032 (ns) | +0.062 |

*Note:* Predicted signs of each discontinuity at the win-loss cutoff. Pure access and the access-dominant allocation share the pattern; the winner and loser magnitudes separate them.

portfolio through the office terms.

The close-election cutoff flips the office indicators of the focal pair while holding everything else fixed. By the dominant-amount pick rule the pick $p$ is more proximate than the rival $r$ ($\rho_p > \rho_r$); as-if-random assignment holds $B$, $V_O$, $V_I$, the proximities, and the rest of the choice set balanced across the threshold, switching only $(o_p, o_r)$. The four outcomes are the resulting discontinuities. Giving to the focal winner and giving to the focal loser are each governed by the proximity gap scaled by $V_I$: the office term enters both limits and cancels, so the discontinuity arises as the label "winner" moves from the proximate pick to the less-proximate rival, not from office-following as such. Both are zero when $V_I = 0$, and they are respectively positive and negative when $V_I > 0$. The change in total giving is zero whenever the budget binds, positive only if $B$ expands on loss. The winnerward partisan shift is governed by $V_O$: reallocating toward whoever holds office tilts the portfolio toward the winner's party, vanishing when $V_O = 0$ and reversing under partisan loyalty.

The mapping turns each estimate into a statement about a primitive. The precise null on total giving (Section 4.1) says the budget binds, rejecting the budget-expansion hybrid.

The positive winnerward discontinuity (Section 4.1) requires $V_O > 0$, rejecting pure ideology and partisan loyalty, under which the portfolio's partisan composition does not track who wins. The same conclusion appears at the candidate level: the committee funds the rival (the opponent it did not previously favor) about ten percent of the time when that rival narrowly wins the focal seat and about three percent when the same rival narrowly loses, a gap that cannot arise under pure ideology because the rival's proximity to the committee is identical either way. Holding office itself carries value. Section 4.1 estimates these candidate-level contrasts directly, as a *continuation premium*, the effect of the pick's win on funding the pick, and a *crossover premium*, the effect of the rival's win on funding the rival. The model signs both positive whenever $V_O > 0$; it does not order their size, because with concave $\psi$ each premium depends on the probability mass near the funding threshold at that proximity level, so the comparison is measured, not derived. The large winner and loser discontinuities establish the converse: their magnitude requires $V_I > 0$. This does not separate access from ideology by itself, but it rules out a knife-edge pure-access account in which proximity has no weight.

Taken together the four outcomes identify $V_O > 0$, $V_I > 0$, and a binding budget, with the relative magnitude $V_O/V_I$ ordered by the substitution test of Section 4.2: giving redirected to the focal winner is more than twice the giving redirected toward ideologically similar non-focal candidates, so the office weight dominates the weight of ideological proximity, a pre-existing relationship, or both, without eliminating it. The object the design identifies is therefore a budget-constrained allocation in which access dominates ideology ($V_O \gg V_I > 0$ with $B$ fixed), not the limiting pure-access case $V_I = 0$, which the winner and loser magnitudes and the documented ideology channels rule out; access-dominant is the precise description.

The partisan-composition outcome does the same work against a partisan-loyalty alternative that the headline winner discontinuity cannot rule out. A loyal committee whose pick loses migrates toward a co-partisan winner wherever one exists, tilting its

portfolio toward its own party, a negative winnerward shift. A positive pooled winnerward shift is therefore the access fingerprint on the partisan-loyalty margin.

The four outcomes adjudicate among the three views but not between two operative sub-mechanisms within the access view. Under a push reading, PACs proactively recalibrate their access portfolios when the office-holding identity changes; under a pull reading, the post-race contributions reflect winner-side solicitation: the new office-holder's staff, party leadership, or fundraising consultants reaching out to the committees that funded the defeated candidate. Contribution data record the transfer and not its initiator, so the two readings imply identical signs on all four outcomes. The readings differ in how redirection responds to winner-side solicitation capacity, and Appendix Section N adjudicates via committee-chair variation.

Two implications follow for heterogeneity. First, ideology predicts attenuation in PAC-by-winner ideological distance, because the policy motive for supporting a distant winner is weak; the access view predicts flatness or amplification, because the value originates in office-holding and information needs are, if anything, greater when the office-holder's preferences differ (the Hall-Deardorff subsidy, which flows to already-aligned legislators, predicts otherwise). Because this moderator conditions on a winner trait, it is exposed to the Marshall (2024) compensating-differentials concern and only partially identified (Section 4.2).

Second, under access the redirection should be flat across the race-level CFscore gap, because PACs seek the office however polarized the race, whereas ideology predicts attenuation in high-gap races. The gap is fixed before the outcome materializes, so this moderator lies largely outside Marshall's objection; the price of that immunity is that it tests the weaker ideology prediction, so I read the two moderators jointly, the race-side test as the cleaner but weaker one and the PAC-side test as the sharper but partially identified one.

A final implication concerns the boundary of the access channel: contributions are

candidate-specific or office-specific, not state-level diffuse transfers. If the redirection were partly a generic state-level access response, giving to non-focal in-state winners would rise on the loss side; the geographic complement I report below tests this prediction directly.

The three views are usually posed as competing accounts of the same donor. They can instead describe different donors. Access-seeking is the logic of investment: the return is legislative attention from whoever holds the seat, so the investor follows the office. Expression is the logic of consumption, support for a candidate or cause (Ansolabehere et al., 2003). An investor reallocates when the office-holding identity changes; a consumer does not, because the object of consumption is the candidate, not the office. Which motive dominates should depend on the donor's position along the spectrum of giving: organized committees (staffed, professionalized repeat players) at the investment end; small individual donors, giving modest amounts to candidates they identify with, at the expressive end; large individual donors, whose giving approaches a committee's scale and repeat structure, in between. Strategic capacity, policy stakes, and the sheer price of access rise together along this spectrum; the design tests the ordering they jointly imply, not their separation.

This reframing yields a prediction the within-committee comparison cannot deliver: the four-outcome signature should appear wherever giving is strategic and fade where it is expressive. Among organized committees, uniformly strategic, the signature should be near-uniform across types (business, trade, labor, ideological, leadership) despite their sharply different professed purposes. Across the full spectrum it should strengthen from the expressive toward the investment end, and the purest access marker, partisan reorientation toward the winning party, which an expressive donor has no reason to perform, should rise from near zero at the expressive end to full magnitude at the strategic end. Every committee is already a strategic donor, so only an extension across the spectrum can test this.

# 3 Research Design

My data come from the Database on Ideology, Money in Politics, and Elections (DIME), assembled and maintained by Bonica (2014, 2018), which records every itemized contribution to every federal candidate from 1980 onward, with donor and recipient identifiers, dollar amounts, dates, and recipient ideal-point estimates (CFscores). DIME classifies each contributor as either an individual (`contributor_type = "I"`) or a committee (`contributor_type = "C"`); the committee category comprises every registered committee filer: corporate, trade, labor, and ideological PACs alongside party and candidate committees. The design runs on both universes giving to U.S. House candidates in general elections, 1990 through 2024: the committee panel, which carries the four-outcome adjudication, and an individual-donor panel built under identical race definitions and pick rules (Section 4.3). Committee coverage is sparse in the 1980s, so the sample begins in 1990; focal cycles run through 2022, with t+2 outcomes measured in the 2024 cycle. DIME's 2024 release omits the recipient election-result fields that define the general-election universe, so 2024 universe membership is reconstructed from MIT Election Lab House returns, a patch that captures 99.65 percent of committee dollars in a 2022 validation (Appendix Section A.1).

The unit of analysis is the committee-by-focal-race observation. For each contested two-candidate U.S. House general election (which I call "the focal race"), I identify every committee that contributed to either of the top two finishers during the focal cycle, and assign each a "pick" using a dominant-amount rule: the candidate to whom the committee gave more dollars over that cycle. The rule yields a unique pick for the great majority of committee-race pairs; the small subset giving to both top finishers is handled under three alternative rules, which move the three probability and composition outcomes by at most one percentage point and the log-giving outcome by 0.011 log points (Appendix Section A.2).

The running variable (rv) is the focal race's vote-share margin, signed positive when

the PAC's pick is the eventual winner and negative when the pick is the eventual loser. The cutoff is zero. The signing convention makes the design a donor-perspective close-election RDD rather than a candidate-perspective one: PACs whose pick barely won (rv > 0) and PACs whose pick barely lost (rv < 0) fall on opposite sides of the threshold, and the discontinuity identifies the behavioral consequence of the focal-race outcome for subsequent giving.

I estimate each discontinuity as a conventional local-linear regression on either side of the cutoff with a triangular kernel and an MSE-optimal bandwidth, and report the robust bias-corrected standard errors and confidence intervals of Calonico et al. (2014). Because the margin varies at the race level, treatment is assigned once per focal race, and every committee that backed either candidate shares that one quasi-random draw; standard errors therefore cluster at the focal race, the level of assignment, which is roughly three times more conservative than PAC-level clustering. Depending on the outcome's bandwidth, the estimates rest on 737 to 1,958 race clusters (Appendix Section A.4). De Magalhães et al. (2025) document nominal coverage for this estimator in close-election RDDs with adaptive bandwidth selection; bandwidths vary by outcome and are reported in Table 3.

My four co-primary outcomes, pre-specified together with the four-model adjudication strategy, are measured at cycle t+2, the next House cycle after the focal race: `gave_to_winner_t2`, an indicator for any positive contribution to the focal-race winner; `gave_to_loser_t2`, its analogue for the loser; `delta_log_total_giving`, the change in the log of the PAC's total House giving (with log(1+x) handling of zeros); and the change in the partisan composition of the PAC's House contributions. I pre-specified the fourth outcome as `delta_dem_share`, the change in the PAC's Democratic share, and report it directly; because its pooled discontinuity is a cancellation artifact of opposite-sign within-pool movements, the primary parametrization is the sign-aligned `delta_toward_winner` (`delta_dem_share` times +1 in D-winning and

-1 in R-winning races), so positive values mean movement toward the winner's party. The family-wise correction applies to the four pre-specified outcomes, and the appendix reports inverse-hyperbolic-sine and unsigned-amount variants of the third. Outcomes are measured at t+2 because t+1 is not well-defined for biennial House cycles and t+4 substantially shrinks the sample.

Two derived outcomes reparametrize the first pair by candidate rather than by office: `gave_to_pick_t2`, an indicator for funding the committee's own pick, and `gave_to_opp_t2`, its analogue for the non-picked focal candidate. Each is a deterministic recombination of the pre-specified pair (on the win side the pick outcome coincides with the winner outcome observation by observation, on the loss side with the loser outcome), so the family-wise correction is unaffected.[1] Their discontinuities are the office premiums of Section 4.1.

The estimand is a between-PAC contrast, not a within-PAC counterfactual: the discontinuity compares the average behavior of PACs whose pick barely won with that of PACs whose pick barely lost, two populations that are statistically equivalent on every observable except the focal-race outcome. The strategic-redirection interpretation rests on the joint structure across the four outcomes rather than on any single discontinuity; any account that fails to match all four signatures is falsified. The within-donor question that the redirection language raises is taken up in Section 4.6, which follows the same PAC's giving to the two fixed focal candidates over time; a within-PAC fixed-effects sensitivity appears in Appendix Section M.

Identification rests on continuity of pre-treatment characteristics through the cutoff (Lee, 2008; Calonico et al., 2014; de la Cuesta and Imai, 2016). Table 2 reports the balance test on six baseline covariates measured before the focal race: all six balance at the cutoff; the largest absolute t-statistic is 1.31, on the number of transactions, and the other five are

[1] The same operation as the sign-aligned `delta_toward_winner` above: a deterministic transform of pre-specified outcomes. The crossover premium is reported loss-side over win-side, so positive values mean the opponent's victory attracts the committee's money.

**Table 2:** Covariate balance at the cutoff.

| Covariate | RD estimate | SE | \|t\| | p-value |
|---|---|---|---|---|
| Log contribution to focal candidate (t) | 0.005 | 0.042 | 0.11 | 0.77 |
| Number of transactions (t) | 0.169 | 0.129 | 1.31 | 0.16 |
| Log total House giving (t) | 0.055 | 0.072 | 0.76 | 0.34 |
| Democratic share (t) | -0.037 | 0.041 | 0.89 | 0.20 |
| PAC age in DIME | 0.173 | 0.339 | 0.51 | 0.65 |
| Focal candidate CFscore | 0.096 | 0.116 | 0.83 | 0.23 |

*Note:* RD discontinuities in pre-treatment covariates: MSE-optimal bandwidth, triangular kernel, race-clustered SEs.

below one.

The McCrary (2008) density test rejects continuity at the cutoff (T = 9.2), with more mass on the winner side. The rejection is mechanical rather than evidence of manipulation: PACs cannot determine close House outcomes, and because winners attract more PAC dollars than losers, the universe of PACs whose pick won is mechanically larger. As de la Cuesta and Imai (2016) emphasize, the appropriate validity check when the running variable is signed by donor-side choice is covariate balance, which the design passes, including inside the inner-donut window ($|rv| < 0.5$ percentage points), where precise manipulation would concentrate. The density diagnostic, bandwidth-stability and donut grids, the inner-donut balance table, and a differential-attrition estimate at t+2 (1.1 percentage points, not distinguishable from zero and too small to move the zero-treated-as-zero coding) appear in Appendix Section J.

Quadratic-polynomial and honest-RD variants are within sampling error of the linear primary (Section 4.7).

Among the heterogeneity tests, the PAC-side ideological-distance moderator condi-

tions on a winner trait and is exposed to the compensating-differentials critique of Marshall (2024); the race-side CFscore-gap and geographic-concentration moderators are predetermined relative to the close-election outcome, so the critique applies to them in much-attenuated form. Section 4.2 develops the division of labor: the cleanly identified race-side test speaks to the weaker ideology prediction, the partially identified PAC-side test to the stronger, read jointly.

Two features of the sample depart from prior close-election work on contributions. First, the committee universe is heterogeneous: business, trade, labor, membership, leadership, and non-connected ideological PACs alongside party and candidate committees, with positively marked corporate committees a small minority. Rather than restrict to one type, I retain the full universe and report the redirection by committee type (Section 4.3); the access-versus-ideology contrast is lost when donors of every motivation enter as one undifferentiated pool, as in the individual-donor design of Dumas and Shohfi (2018). Second, I build a parallel panel of individual donors under identical race definitions and pick rule (20.4 million donor-by-race observations across 6.3 million donors, focal cycles 1990 through 2022), so the two designs differ only in the donor universe. The committee sample (1,806,639 committee-by-race observations, 5,969 contested races, 117,251 unique committees) is to my knowledge the largest behavioral close-election RDD on organized giving.

# 4 Results

## 4.1 The four-outcome joint signature

At the cutoff, committee behavior two years after the focal race reduces to four rates: committees whose pick narrowly won fund that candidate, now the officeholder, 48 percent of the time and the defeated opponent 3 percent; committees whose pick narrowly lost fund their defeated pick 26 percent of the time and the winner who beat them 10

percent (side-specific limits at each outcome's bandwidth, 15.1 and 6.6 pp; Appendix Section B.4). Reading the two-by-two by candidate gives the candidate-level estimands of Section 3. The *continuation premium*, the effect of the pick's win on funding that same pick, is +20.0 percentage points (Table 3; SE 3.1); the *crossover premium*, the effect of the opponent's win on funding that opponent, is +6.9 points (SE 1.0): committees become measurably likelier to fund the other side's candidate for no reason other than that the candidate now holds the office. Both premiums are positive, the scale-invariant part of the result, which no expressive account delivers. Their relative size is not scale-invariant (continuation larger in percentage points, crossover larger against its 3-percent baseline), so I read the asymmetry as description, not mechanism.

Read by office, the same levels compound into the composites the close-election literature reports, retained for comparability. The probability of giving to the race winner jumps 37.5 percentage points (Table 3; SE 1.6 pp; 95% CI [0.349, 0.411]; bandwidth 15.1 pp; effective N approximately 710,000), a contrast of levels dominated by win-side continuation, not a redirection rate. These are large effects by the standard of the committee-access literature: exile from a major money committee costs a legislator on the order of $100,000–$150,000 in giving from the industries under its jurisdiction (Powell and Grimmer, 2016), and joining a committee raises jurisdiction-relevant giving by roughly 27 percent (Fouirnaies and Hall, 2018); the margin here is prior to every committee threshold, whether the candidate holds office at all. The effect rises monotonically in PAC size (small +0.27, medium +0.40, large +0.48; Appendix Section B.2), so no handful of mega-PACs drives it.

The substitution leg is just as sharp. The probability that the PAC continues giving to the focal-race loser drops by 22.7 percentage points (Table 3, column 2; SE 2.3 pp; CI [-0.284, -0.192]; bandwidth 6.6 pp): defeat nearly halves loyalty to the pick. What the defeated candidate does next decides what the drop means, and candidacy is measured at three tiers that differ by a factor of two: at the cutoff, 88 percent of narrow losers file with the FEC for some office at t+2 and 82 percent for the same seat, but only 32 percent are on

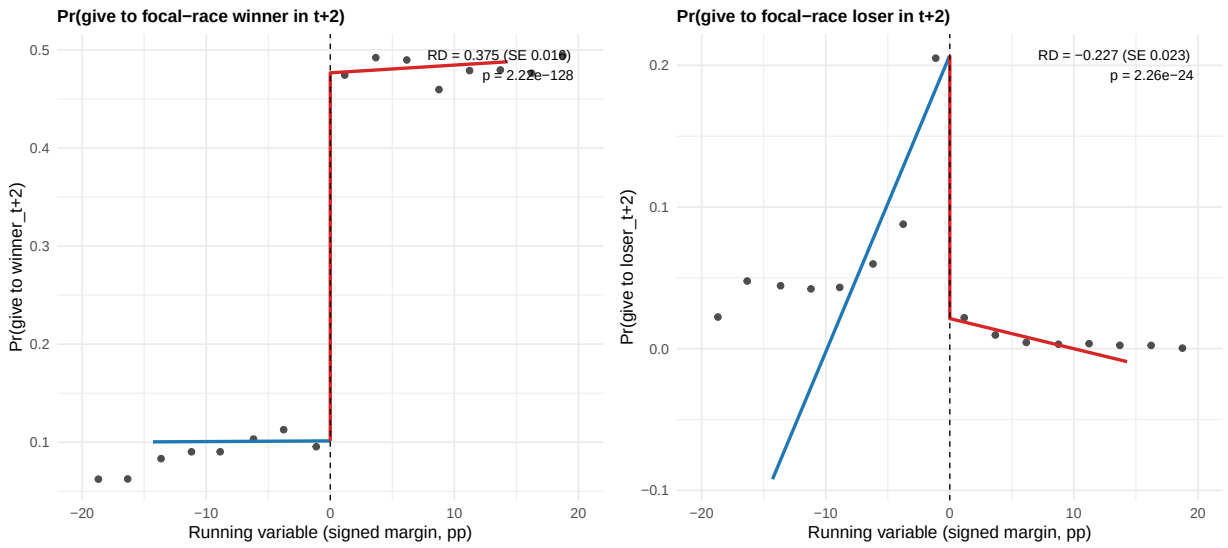


**Figure 1:** Probability of giving to the focal winner (left) and loser (right) in the next cycle, at the cutoff.

the same seat's general-election ballot again (Appendix Section L). Splitting the design on that ballot-verified rematch is decisive. Where the defeated candidate runs again, the candidate-level premiums vanish: backers of the narrow loser fund that candidate 44 percent of the time, the same rate at which backers of the narrow winner fund the winner (a continuation premium of +0.8 points, SE 4.4), and neither side's backers cross to the other candidate (5 percent each). Where the defeated candidate has left the ballot, 8 percent of the committees that backed the loser still fund that candidate, 13 percent fund the winner they had opposed, and the two premiums are +40 and +12 points. The office premiums of Table 3 are the premiums of staying in contention for the seat: office keeps a committee's relationship alive by keeping the candidate in the field, and once the defeated candidate leaves the field the committee moves to the officeholder rather than to nobody. The split conditions on the loser's decision to run again, made after the outcome, so the rematch rates bound conditional loyalty from above (Appendix Section L).

The hybrid account predicts that the redirection is additive: winner support added without loser support subtracted, expanding total giving. The data show no expansion. The change in log total House giving across the cutoff is +0.032 log units (Table 3, col-

umn 3; SE 0.131), indistinguishable from zero. Its confidence interval is nonetheless wide (the upper bound corresponds to roughly a 31 percent expansion), but the width is an artifact of the logarithm amplifying committee exit: roughly one in seven committees ceases House giving by the following cycle, and exit is the opposite of what a hybrid requires. Measured on the margin the hybrid actually predicts, the surviving committee's spending envelope, the budget is tightly fixed: a loss changes total giving by +0.6 percent under within-committee fixed effects, +1.9 percent among committees that remain active, and +0.1 percent as a share of the prior budget, bounding expansion below 5.1 to 7.3 percent on these exit-robust margins while the unconditional log-change remains the wide $\pm$31 percent (Appendix Section Q). The redirection is financed by reallocation within a near-fixed budget, the budget-constrained allocation of the Section 2 model, in which access competes with loyalty for a fixed envelope.

The fourth co-primary outcome, the winnerward shift in the PAC's partisan composition, is a positive 6.23-percentage-point discontinuity (Table 3, column 4; SE 0.80 pp; bandwidth 12.3 pp).[2] The pre-specified raw outcome, the change in the PAC's Democratic share, is essentially zero in pool (Table 3; +0.010, not distinguishable from zero): a cancellation artifact of opposite-sign within-pool movements, not a substantive null on partisan loyalty. Because the outcome is computed on the PAC's entire House portfolio, its anatomy matters. Removing the focal-race contributions leaves +4.97 percentage points, but most of that remainder is the committee's redirection in its other contested races, projected onto the focal winner's party through its same-party slate (an accounting property of the measure for multi-race committees, not a causal response). Excluding every race where the committee holds a pick leaves +0.9 points pooled (not distinguishable

[2]I report this outcome with the per-observation Democratic-share change multiplied by an indicator that takes value +1 in races won by a Democrat and -1 in races won by a Republican. Positive values therefore indicate movement toward the focal-race winner's party in every pool. The same convention is used for the robustness grid (Appendix Table B1: logit RDD and race-margin LPM). The rdrobust coefficient on the sign-aligned outcome is reported as the absolute magnitude in the substantive prose (loser-side limit minus winner-side limit) so that positive values represent winnerward redirection by loser-PACs over the winner-PAC baseline.

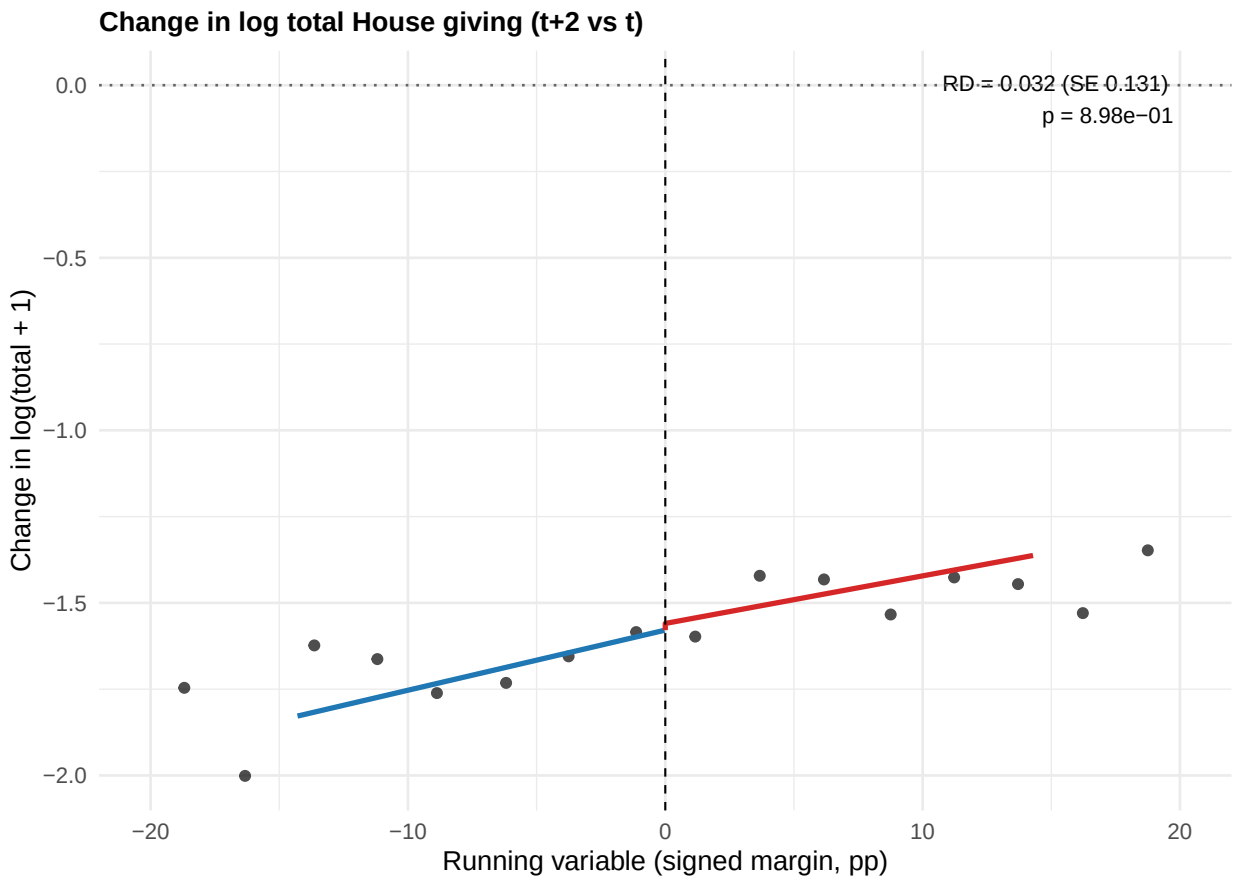


**Figure 2:** Change in log total House giving at the cutoff: no discontinuity.

from zero) and +3.0 points after incumbent defeats (more than three standard errors from zero; Appendix Section O).

A second discipline on the estimate is a pre-period differential trend that the design itself predicts. The running variable is the committee's own signed margin, so within a single close race the winner's backers sit just above the cutoff and the loser's backers just below; close-election assignment randomizes which side won, not which committees backed whom. The two groups lean toward opposite parties, and committee giving sorted toward its own party over these decades, so the winner-oriented share was already widening across the cutoff before the focal race: by 1.65 percentage points (SE 0.67) over the two cycles preceding it, and by 2.33 points (SE 0.66) over the two cycles before that, measured with the focal candidates' contributions removed from the share

at every endpoint (Appendix Section J). A difference-in-differences adjustment subtracts each committee's pre-period drift from its post-period shift, under the assumption that the differential trend would have continued at its measured rate, an extrapolation the two independent pre-period windows support. The trend-adjusted, focal-excluded shift is +6.81 percentage points (SE 1.00); the focal-inclusive counterpart is +8.67 points (Appendix Section J). Before the race, the losing side's portfolio was moving away from the eventual winner's party relative to the winning side's; after the race it moves toward it. The composition outcome tells against partisan loyalty through the direction of that reversal (toward the winning party, where loyalty predicts continued movement away), not through portfolio-wide reorientation.

The shift decomposes symmetrically across winner parties. On the same focal-excluded, trend-adjusted basis, PACs move toward the winning Democrat when a Democrat wins (+4.3 percentage points, SE 1.5) and toward the winning Republican when a Republican wins (+8.2 points, SE 1.3; Appendix Section J). Co-partisan migration is incompatible with both pools, and on the raw Democratic-share scale the two movements (+5.4 and $-6.6$ points) produce the pooled cancellation.

After Bonferroni and Benjamini-Hochberg adjustment for the pre-specified four-outcome family, the two large discontinuities survive with adjusted p-values below 10^-20 and the two nulls remain non-significant by construction (Appendix Section K).

The four outcomes form a single signature rather than four separate findings, and the model of Section 2 assigns the adjudication load unevenly. A pure-ideology account (including its partisan-loyalty variant) is falsified by the winnerward partisan-composition discontinuity together with the crossover premium: both require office-holding to carry value ($V_O > 0$), which ideology sets to zero, and partisan loyalty in particular predicts a negative or zero winnerward shift where the data show a sharply positive one. The large winner and loser discontinuities do not by themselves falsify ideology, since ideology generates the same signs through relabeling (Section 2), but their magnitude rules

out the opposite knife-edge, a pure-access account in which proximity has no weight. A budget-expansion hybrid is falsified by the null on total giving, near-fixed on precisely the margins expansion would move. An information account, in which losing reveals the candidate is unlikely to recover, predicts a smaller effect in open-seat races; the open-seat effect, +0.43, is indistinguishable from the +0.44 for re-elected incumbents, a difference of 0.007 with a standard error of 0.028 (era-adjusted +0.43 against +0.44; Appendix Section B.2), so the prediction finds no support. Access is the only account left standing.

The crossover premium also tracks how the seat changed hands: +9.6 percentage points when the committee's defeated pick was the incumbent (so the winner is new to the seat), +8.4 in open races, and zero (+1.2, indistinguishable from zero) when the pick was a challenger and the opponent a sitting incumbent who survived (Appendix Section B.4). Each estimate is causal within its configuration; comparison across them is descriptive, since committees choose whom to back. With that caveat, the gradient is the one access predicts: committees cross over when office actually changes hands, and not at all toward an incumbent they had funded a challenger against.

## 4.2 Access versus ideology

The signature leaves one serious rival standing: ideology. Three tests within the same quasi-experiment adjudicate between them: a direct measurement of where the loser's dollars would otherwise go, a race-side moderator that is cleanly identified but tests the weaker ideology prediction, and a PAC-side moderator that tests the stronger prediction but is only partially identified. I take the three in turn and then read them jointly.

The direct measurement targets the default-to-winner reading, under which the PAC redirects to the focal winner only because no ideologically similar alternative exists in its choice set. I track PAC giving to non-focal House candidates whose CFscore lies within 1.38 units of the focal loser's (approximately the 25th percentile of the race-level winner-loser gap, tight enough that the binary outcome does not saturate). The default reading

fails on both of its predictions. Replacement candidates are widely available: at baseline, more than three-quarters of PACs give to someone within ideological proximity of the focal loser. And loser-PACs do use them: giving to an ideologically similar non-focal candidate rises 16.1 percentage points at the cutoff (SE 1.7 pp), yet the focal winner still receives 2.3 times that redirection. Looser thresholds, an intensive-margin count, and the full sensitivity curve give the same result (Appendix Section H); and since noisy CFscores attenuate the estimate (Bonica, 2014), doubling it under a generous reliability ratio narrows but does not close the gap with the +0.375 focal channel. What the secondary channel *is* remains ambiguous (ideology or an existing relationship, which are observationally equivalent here), so I read this channel as secondary substitution within the PAC's relationship space rather than an identified ideology channel, a distinction that leaves the access interpretation intact either way.

A geographic complement pins the channel to the candidate rather than the state. Tracking giving to non-focal winners in the focal state returns +0.011 (SE 2.2 pp), indistinguishable from zero; the design detects only shifts above about four percentage points, well below the 16.1-point rise toward ideologically similar candidates. PACs do not redirect toward generic in-state winners but toward the specific person who replaced their pick, in line with the relationship-specific access view that goes back to Welch (1980) and recently surveyed by Phillips (2024); Hall and Deardorff (2006) supply the closely related lobbying-side framing.

The race-side moderator conditions on the ideological gap between the focal winner and loser, the absolute CFscore distance fixed before the close outcome materializes. That timing makes it largely immune to the Marshall (2024) concern developed below; the price of that immunity is that it tests the weaker ideology prediction, the race's polarization rather than the PAC's own ideological cost. Pure ideology predicts attenuation in the gap, since crossing a wider distance is harder; access predicts flatness. The data show flatness: redirection across gap quartiles is +0.395, +0.408, +0.380, and +0.356 (a four-point

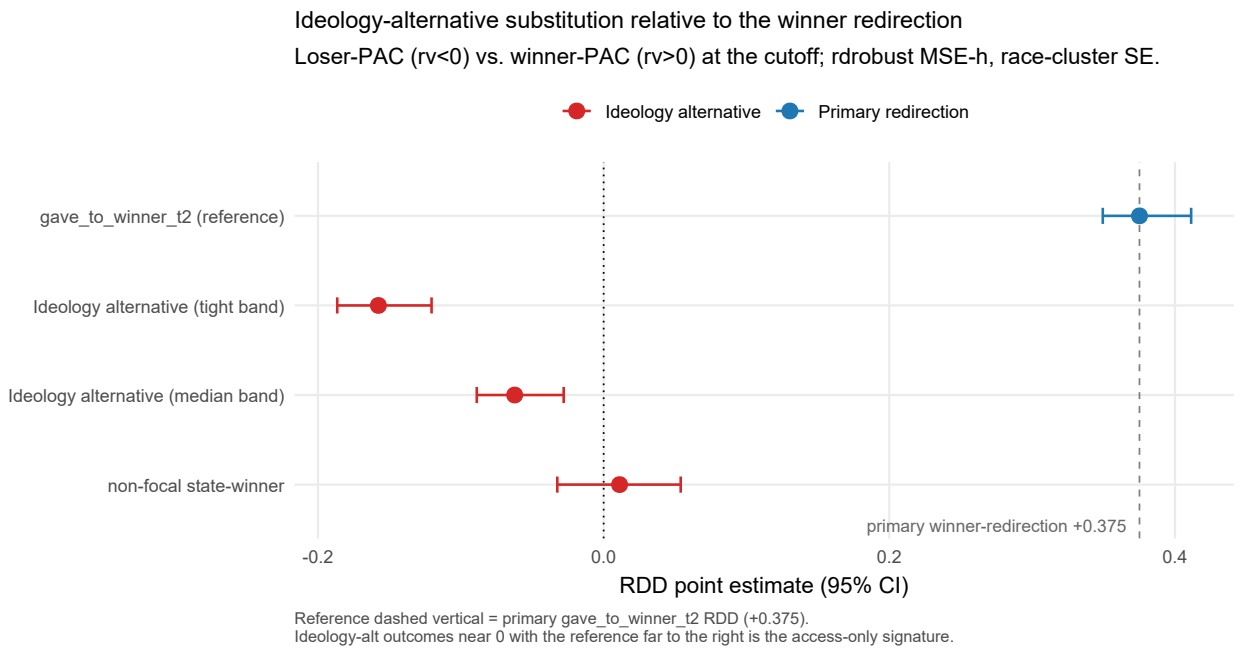


**Figure 3:** Ideology-alternative substitution outcomes compared with the primary winner redirection.

Q1-to-Q4 difference on a 40-point base), and the continuous loss-by-gap interaction is -0.019 (SE 0.014; Appendix Section G), indistinguishable from zero. Even doubling that slope under errors-in-variables in the moderator (Bonica, 2014) leaves it below any plausible ideology-loaded benchmark. The cleanest test the design offers finds no ideology gradient on the prediction it can reach.

The PAC-side moderator conditions on the CFscore distance between the contributing PAC and the focal-race winner, and tests the substantively stronger prediction: whether the PAC's own ideological cost of funding a distant winner moderates its redirection. This is the margin where access and ideology compete most directly, but also where identification is weakest. Because the moderator is a function of a winner trait, it inherits Marshall's (2024) compensating-differentials problem: candidates who win close races despite an unfavorable observed trait must be advantaged on unobservables, so conditioning on the trait reintroduces exactly the selection the design removes. The interaction magnitude is therefore not point-identified. Descriptively, redirection declines monotoni-

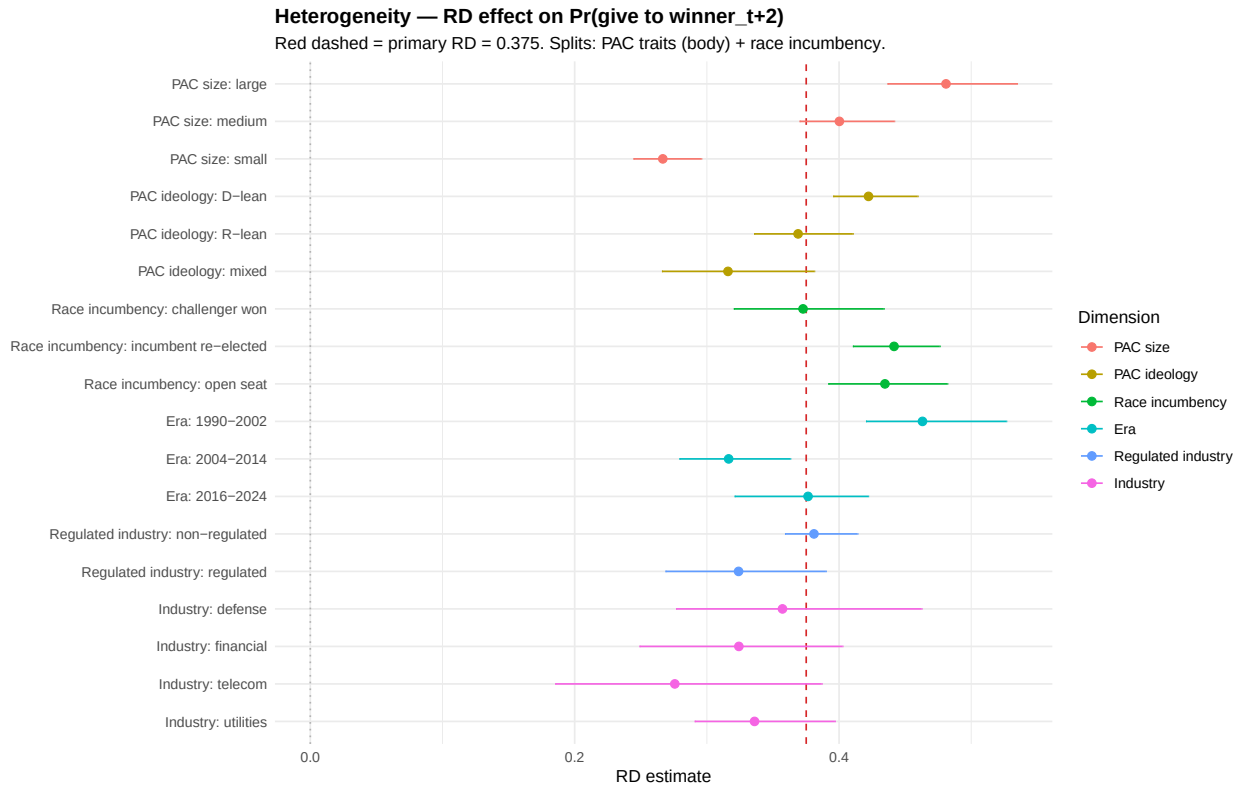


**Figure 4:** Winner redirection by PAC size, PAC ideology, race incumbency type, era, and industry.

cally in distance: same-tercile PAC-winner pairs +0.368, adjacent-tercile +0.330, opposite-tercile +0.211, with a continuous interaction of -0.226 (SE 0.013) at face value. But applying the partial-identification bounds of Marshall's section 6 with the local CFscore variance $\sigma$_Z $\approx$ 0.92 produces intervals that bracket zero at every considered tightness (K = 0.5: [-0.69, +0.23]; Appendix Section F). I therefore retain the tercile pattern as descriptive evidence of a layered ideology channel running alongside the dominant access pattern, without interpreting its magnitude.

The three tests are complementary rather than convergent, and none is individually decisive. The race-side test cleanly rejects the weaker ideology prediction; the PAC-side test finds an ideology gradient real in sign but bounded away from any magnitude large enough to overturn access; and giving redirected to the focal winner is more than twice the quantified secondary substitution channel. Read together they support the access-dominant reading ($V_O \gg V_I > 0$, not the knife-edge $V_I = 0$), with the headline +0.375

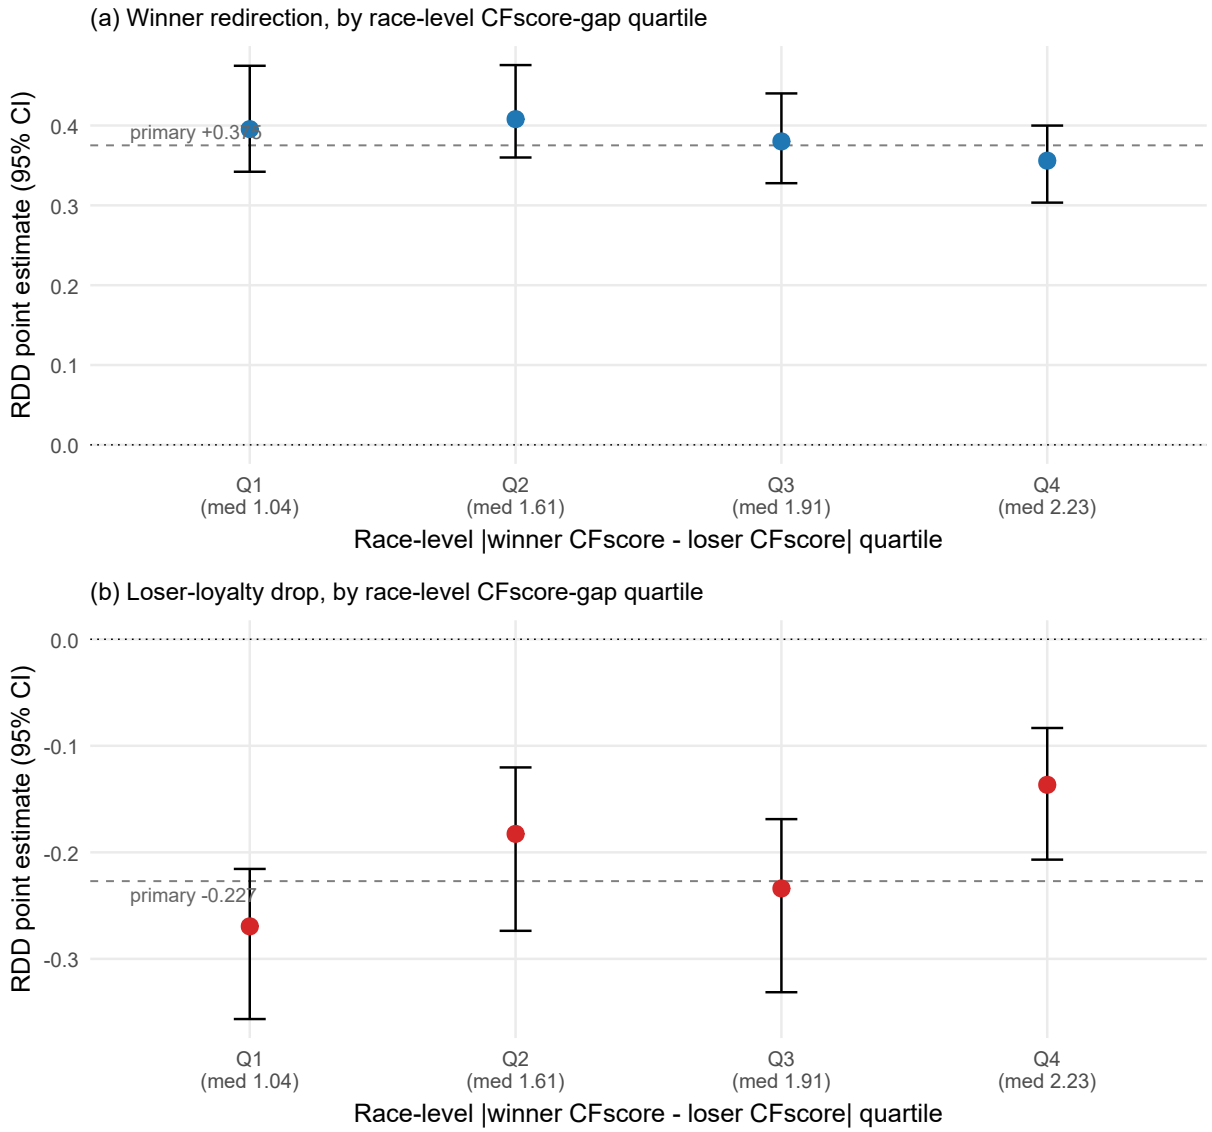


**Figure 5:** Winner redirection (top) and loser-loyalty drop (bottom) by race-level CFscore-gap quartile.

unaffected by any of the three. Loser-loyalty, the one non-flat race-side outcome, weakens slightly less in high-gap races (interaction -0.073, significant at the 0.1 percent level); the interaction preserves the sign and does not bear on the winner-side adjudication. An industry split makes the same point about generality: utilities +0.34, financial services +0.32, telecommunications +0.28, defense +0.36, non-regulated +0.38 (Appendix Section

B.2); no single sector drives the headline, and the pattern does not concentrate in regulated sectors.

## 4.3 Strategic versus expressive donors

Section 2 implies a test the committee data can run only halfway and the individual data can complete: the signature should be near-uniform among organized committees and scale down toward small expressive donors.

I split the committee sample by donor type (Table 4; construction in Appendix Section P). Every organized class redirects strongly (business-connected +0.374, labor +0.367, non-connected +0.456); the classified-sample re-run reproduces the headline (+0.374), and excluding party and candidate committees leaves +0.357. Because non-connected status is an FEC filing category rather than a measure of purpose, I classify purpose directly with the Center for Responsive Politics taxonomy: ideological and single-issue committees, the donor type an expressive account would single out as least transactional, redirect among the most (+0.451 and +0.494) and convert to former opponents besides (+0.081 and +0.044), differences traceable to baseline portfolio structure, not professed purpose: the crossover premium is paid by the committees whose giving already spans both parties, and within that group it is a jump at the cutoff rather than a standing two-party allocation (Appendix Section P). The type split delivers the first half of the prediction (strong wherever giving is organized) but not the second, because every committee is already strategic.

To complete the test I run the identical design on individual donors, who span the full donor spectrum: 20.4 million donor-by-race observations on the committee panel's races and focal cycles, stratified at fixed cutoffs into small (under $200), mid, and large (over $500) donors (Appendix Section P). Individual donors as a whole redirect toward the winner less than committees do (+0.201 versus +0.375); the gradient is the test: redirection rises monotonically with donation size, +0.144, +0.300, and +0.472 across the strata.

Because the binary outcome is partly mechanical (a donor who gives to more candidates is more likely to give to any one of them), I read the gradient primarily off the partisan-reorientation share, which is invariant to giving breadth and which an expressive donor has no reason to move. The break falls between small donors and everyone else: 0.005, an order of magnitude below the committee 0.062, against 0.042 among mid-sized and large donors alike. The gradient is behavior, not breadth: small donors are more active at t+2 than mid-sized donors, and conditioning on t+2 activity leaves the gradient intact and sharper (+0.25, +0.56, +0.65; Appendix Section P).

The fourth co-primary outcome separates the two donor classes. Committees redirect within near-fixed budgets (the total-giving null of Table 3), while individual giving is itself outcome-contingent: total House giving jumps +0.48 log points among large donors (SE 0.15, bandwidth-stable) and +0.33 among mid-sized (SE 0.14, uncorrected), through both margins, since giving at all next cycle rises 5.4 percentage points and totals among donors active in both cycles rise roughly 11 percent. Because FEC rules mandate itemization only above $200 per recipient and cycle, small-donor totals are incompletely observed; the budget test is read on the fully itemized large stratum (Appendix Section P).

The premium decomposition of Section 4.1 completes the anatomy. Large donors' own-pick premium matches the committee value almost exactly (+0.21 versus +0.20); the crossover premium does not carry over: +0.011 among large donors, +0.006 among mid-sized, zero among small, against +0.069 for committees. Funding the opponent who beat one's candidate, the component ideology cannot generate, remains an organizational behavior. The composition of the strata matches this reading: among close-race donors in the 2020–2024 files, policy-exposed occupations (executives and owners, attorneys, finance and real estate, physicians) are roughly a third of large donors against an eighth of small (Appendix Section P).

None of this is an artifact of the small-donor era: the gradient and the large-donor budget response reproduce on focal cycles 2020–2022, where four-fifths of donor-race

observations come from donors giving under $100 (Appendix Section P). The gradient relocates the access-versus-expression debate: the operative contrast is not between business and ideological donors but between strategic and expressive giving. Large individual donors share the strategic margins (retention at the committee rate, realignment, outcome-contingent budgets); conversion is the one margin they show only faintly. Small donors express; large donors invest in their own winners; organizations alone routinely buy access from former opponents.

## 4.4 A second channel: directed lobbying follows the winner

The evidence to this point is a pattern in *giving*. If it reflects access-seeking, the same redirection should appear in a second, behaviorally independent channel: lobbying, the activity through which the office term $V_O$ of Section 2 is actually exercised. Both channels are observable together on the subsample of committees sponsored by publicly traded firms, whose lobbying can be matched to specific legislative targets through the firm-by-legislator panel assembled by Kim et al. (2026) (focal cycles 2008 through 2016); I rerun the close-election design with the firm's directed lobbying of the race winner as the outcome.

The discontinuity reproduces. A firm whose backed candidate barely won directs significantly more lobbying expenditure toward that winner over the following cycle than a firm whose candidate barely lost: the local-linear discontinuity is +1.24 log points in directed lobbying dollars (robust SE 0.45), stable across bandwidths from 5 to 20 percentage points. The same firms' donations, measured in the same data, jump in parallel (binary +0.17, log-dollar +1.41, both significant at the 0.1 percent level), so the lobbying result behaves like the giving headline, not an idiosyncratic subsample. Lobbying directed at the *loser*, who is out of office at t+2, shows no discontinuity (-0.29, indistinguishable from zero against the +1.24 for the winner): the effect is specific to the candidate who actually holds the seat, as an access account predicts and a partisan-loyalty account does not.

The result survives the probes a skeptic would run first (Appendix Section R): it holds when the outcome window is restricted to the post-election year alone; it is marginally larger among firms that lobby at all (+1.42, significant at the 1 percent level), so it is a discontinuity in *how much* a firm lobbies the winner, not in *whether* it lobbies; directed lobbying of the committees on which the winner sits jumps as well (+0.15, significant at the 1 percent level); firm-side predetermined covariates are balanced at the cutoff; and the race-level margin density is continuous (the density test does not reject) once the mechanical firm-count replication is removed.

This is mechanism confirmation on the subsample where both channels are observable, not a claim about the full 1990–2024 universe: the firms here are public-company committee sponsors, the window is 2008 through 2016, and the lobbying magnitude is not the giving magnitude. Within those bounds the redirection is no longer only a behavioral signature consistent with access: a parallel, winner-specific discontinuity appears in the activity through which access is exercised, which is what distinguishes access-seeking from the alternatives the giving pattern alone could not separate.

## 4.5 Era heterogeneity

The access redirection has weakened over time but remains substantial. Splitting the sample at policy-relevant cutpoints yields subsample estimates of +0.463 (1990–2002, pre-BCRA), +0.316 (2004–2014), and +0.377 (2016–2024, whose focal cycles run 2016–2022; Appendix Table B1). The formal era-by-loss interaction rejects equal effects across eras at any conventional level (Wald F of 13.94), and every era's effect remains large and significantly positive: the institutional regime change reduces but does not eliminate access-seeking redirection. Nor is the decline a compositional artifact of DIME's coverage drift: the 1,262 committees observed across all three eras show essentially the same decline pattern at a higher level, a separate finding about long-lived committees rather than a threat to the era pattern (Appendix Section E, Table E1).

Two readings are consistent with the aggregate pattern: post-2002 institutional changes (the soft-money ban, then 501(c)(4) and independent-expenditure growth) may have shifted marginal access purchasing into vehicles direct PAC giving does not capture, or growing polarization changed the propensity to redirect. The within-vs-cross-party era split (Appendix Section E) decomposes the attenuation: within-party redirection drops sharply from +0.401 in 1990–2002 to +0.174 in 2004–2014, recovering partially to +0.267, while cross-party redirection is essentially stable at +0.483, +0.430, and +0.428. The attenuation is overwhelmingly a within-party phenomenon, which rules out the polarization story in which cross-aisle support becomes less palatable, since that version predicts the cross-party component to shrink most and it is the stable one. A version in which polarization tightens within-party loyalty fits the data and is observationally equivalent here to channel substitution; my design does not separate the two.

The attenuation parallels Myers (2026), who reports a roughly 50 percent decline in donor punishment of extremism since 2000: donor responsiveness to electoral signals appears to have muted across mechanisms in the post-2002 environment.

The cross-party component, the stable half of the era pattern, is also the half that most sharply identifies the mechanism. Pooled across the three eras, committees whose ideological lean is opposite the focal winner's party redirect at +0.452, against +0.276 for the within-party component (difference z = 4.3). In the allocation model of Section 2, ideological proximity pulls a committee away from a cross-party winner, so a positive cross-party discontinuity isolates the office term $V_O$: the behavior of a contributor that values incumbency enough to fund a legislator it is predisposed to oppose. Decomposing that component shows the office-seeking is broad-based. Committees in classically regulated industries (utilities, finance, telecommunications, and defense) redirect at +0.381, statistically indistinguishable from the +0.456 of the rest of the committee universe (z = -1.6): cross-party office-following is a general property of organized giving, not a behavior concentrated among the committees whose dependence on government is most conspicuous.

Nor does geography concentrate it: the cross-party concentration gradient does not reach conventional significance (Appendix Section E.2), full-sample per-tercile discontinuities are statistically indistinguishable (+0.37, +0.40, +0.36), and the robust continuous interaction implies a slightly *smaller* discontinuity at high concentration, reflecting the narrower portfolios of locally concentrated committees rather than district-specific amplification (Appendix Section I).

The sharpest evidence is in ideological space. Were the cross-party effect an artifact of committees whose ideological sign is barely defined, it would vanish among contributors unambiguously opposed to the winner. It does not: sorting cross-party dyads by the ideological distance between committee and winner, redirection is +0.38 for the closest third and +0.34 for the most distant, a third that funds winners roughly 1.8 CFscore units away, and the difference between the two (0.05, standard error 0.06) is not distinguishable from zero. Office-seeking overcomes even a large ideological cost, and on this margin the cost does not register at all; the proximity weight the model retains is identified elsewhere, in the substitution test of Section 4.2 and in the one-party portfolios of Section 4.3. This is the access-dominant region of the Section 2 model ($V_O$ dominating but not extinguishing $V_I$), and a robust positive effect among the most opposed dyads is one ideology alone cannot generate.

## 4.6 Within-PAC dynamics

The discontinuity estimates compare different committees on either side of the threshold in a single post-period. To ask whether the redirection is a behavioral response *inside* a given committee, I track each committee's contributions to the two fixed focal candidates (the eventual winner and the eventual loser) across the two cycles before and after the focal race, restricting to close contests in focal cycles through 2020, where the full event window exists; Appendix Section M.1 reports the specification and full estimates, and Figure 6 plots the trajectories. Among committees whose pick narrowly lost, giving to

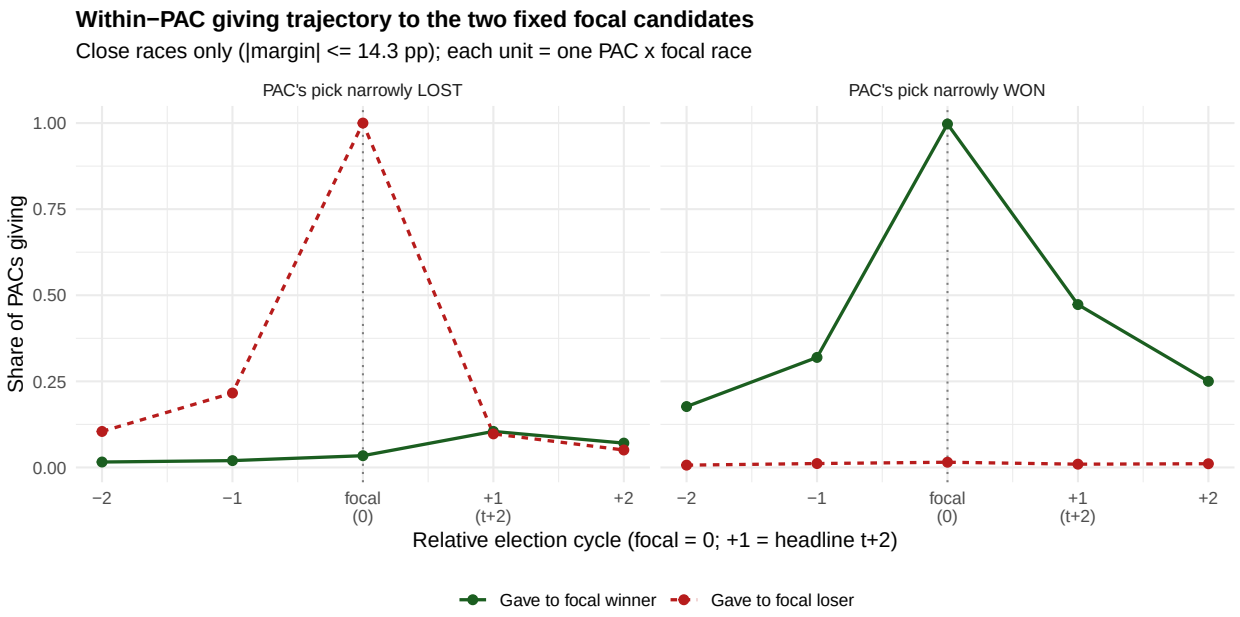


**Figure 6:** Within-PAC giving to the two fixed focal candidates across relative cycles -2 to +2, by focal outcome (+1 = the t+2 outcome cycle).

the winner (the opponent who defeated the committee's pick) is essentially absent before the focal race (about two percent, flat across both pre-cycles) and rises to roughly ten percent two years after the loss, a within-committee increase of eight to nine percentage points (significant at any conventional level) with no detectable pre-trend. Giving to the former pick peaks in the focal cycle and then falls to and below its earlier baseline, so the fall in loser-loyalty is in part the mechanical end of the campaign cycle, and the winner-following rise is the clean behavioral signal. The within-committee preference between the two candidates flips accordingly, a swing of about ten percentage points against the pre-campaign baseline (significant at any conventional level): the redirection is visible as a within-donor reorientation, not only as a cross-sectional discontinuity.

The trajectories reproduce Section 4.1's decomposition inside donors: the win-side level is continuation with the office-holder the committee helped elect, and the within-committee switch toward a former opponent is the crossover premium realized donor by donor.

## 4.7 Robustness

The redirection is stable under the estimator perturbations a technical referee runs first: bandwidths from one to twenty percentage points, quadratic and linear local polynomials, donut-hole exclusions from 0.25 to 2 percentage points, PAC-level rather than race-level clustering, and covariate adjustment (+0.372 against +0.375 unadjusted); the honest-RD inference of Armstrong and Kolesár (2018) returns +0.390 at smoothness bound M = 0.1 and +0.392 at M = 0.5. Appendix Table B1 reports the full grid, Appendix Section B the covariate-adjusted and parametric variants.

More demanding than tuning one estimator are the choices that define the design itself, and the discontinuity survives each of them (Appendix Table B3). It is not an artifact of the extensive-margin indicator: dollars to the winner jump at the threshold, and the winner's share of the committee's post-election House giving (a compositional measure that cannot rise merely because a committee funds more candidates) increases by three percentage points. It reproduces exactly when the panel is collapsed to one observation per race, ruling out any artifact of races with many committees. It is essentially invariant to the pick rule, which can disagree only for the roughly one percent of committees splitting giving between both focal candidates. It is positive and significant in every subsample, smaller though still economically large for single-close-race committees, in line with the strategic-donor gradient of Section 4.3. The one specification whose confidence interval essentially does not overlap the primary is the pre-BCRA subsample, the substantive heterogeneity of Section 4.5 rather than a robustness concern.

A final sensitivity absorbs all time-invariant committee heterogeneity: a parametric local-linear specification with PAC and cycle fixed effects (identification restricted to committees observed under both win and loss states within the bandwidth) yields -0.392 (SE 0.011) on the binary headline and +0.061 (SE 0.006) on the sign-aligned composition derivative, within 1.06 and 0.11 standard errors of the between-PAC estimates (Appendix Section M): removing between-PAC composition does not attenuate the discontinuity.

The contrast is a sub-population sensitivity, identified off the roughly 30 percent of committees observed under both outcomes; the direct within-donor evidence is Section 4.6's event study.

# 5 Conclusion

Three decades of debate over what political money buys have been framed, implicitly, as a question about donor *type*: whether corporations give transactionally while ideological groups give expressively. The evidence here relocates the question. Winner-following is not a property of corporate status but of strategic giving. When a committee's pick narrowly loses, committees of every stripe follow the winner: business, trade, labor, leadership, party, and, as strongly as any, the ideological and single-issue committees an expressive account calls least transactional. Individual donors split along the same line by the scale of their giving: small donors follow only weakly, barely realign, and (in their itemized giving) give no more after a win than a loss, while the largest donors retain their winners at the committee rate, realign toward the winning party, and are mobilized by victory itself. What no individual stratum reproduces is the organizational signature's conversion component, funding the opponent who won. The access-versus-expression contrast is real, but it is a contrast in the scale and organization of giving rather than in donor type, which is why single-type studies have read the same behavior in conflicting ways.

For organized committees, the predicted behavior is within-budget office-following. A committee whose candidate loses does not spend more; it reroutes a near-fixed budget off the defeated candidate and onto the winner, and it tilts the partisan composition of its whole portfolio toward whichever party now holds the seat. The price of office is paid on both relationships: a continuation premium along the committee's own tie and a crossover premium across it, both paid where the seat changes hands and the defeated

candidate leaves the field; where the loser contests the seat again, committees on both sides stay with their candidate. On the public-firm subsample, where lobbying can be matched to legislators, the redirection is accompanied by a parallel, winner-specific jump in directed lobbying, the activity through which access is exercised. That corroboration is scoped to where both channels are observable; on the full universe the access reading rests on the joint signature, consistent with Fowler, Garro, and Spenkuch's (2020) tightly bounded null on firm value: giving reacts to the office flip while firm value does not, a subsidy paid for access to a process rather than for a cash-back outcome.

The reading has normative implications the access-versus-ideology debate has only partly registered. If organized money follows office-holding rather than party, reform arguments premised on money being ideologically aligned with one side miss the system's central tendency: committees back the side that holds power, symmetrically in Democratic- and Republican-won races alike. The mechanism is access-dominant rather than pure access: a bounded ideology channel operates alongside it and survives even large committee-winner ideological distance.

Two open questions bound the claim. The first is where the money went: the direct-PAC access channel has attenuated by roughly a fifth since the Bipartisan Campaign Reform Act, and my design cannot separate access flows migrating into post-Citizens-United vehicles from polarization tightening within-party loyalty; a parallel close-election design on independent-expenditure flows is the natural next step. The second is who initiates the redirection. The design records the transfer but not its opener, and two indirect tests point away from winner-side solicitation: open-seat winners, who arrive with no donor-relations apparatus, generate redirection indistinguishable from re-elected incumbents', and the winner's standing-committee chairmanship returns a null excluding amplification larger than about half the headline. Whether the relationship is opened by the committee or by the winner's office, and what the purchased access consists of, are distinctions contribution data cannot draw.

The right description of strategic political giving in the United States is that it follows the winner. The committees that moved from Pete Gallego to Will Hurd were not converted to Hurd's politics; they were staying with Texas's Twenty-Third District seat. Contributions follow whichever office can advance the donor's agenda, wherever giving is an investment in access rather than an act of expression, and the close-election design that isolates the office-versus-no-office margin is general enough to support the next generation of work on access purchase in the post-Citizens United environment.

**Table 3:** Main RD estimates.

| Outcome (cycle t+2) | Estimate | SE | p-value | h (pp) | Eff. N |
|---|---|---|---|---|---|
| Gave to focal winner | 0.375 | 0.016 | < 10^-127 | 15.1 | 710,846 |
| Gave to focal loser | -0.227 | 0.023 | < 10^-23 | 6.6 | 336,035 |
| Funded own pick (continuation premium) | 0.200 | 0.031 | < 10^-8 | 6.5 | 331,423 |
| Funded non-picked candidate (crossover premium) | 0.069 | 0.010 | < 10^-9 | 9.0 | 434,379 |
| Change in log total House giving | 0.032 | 0.131 | 0.90 | 15.7 | 735,469 |
| Change in Democratic share (pre-specified) | 0.010 | 0.009 | 0.14 | 11.4 | 466,924 |
| Winnerward shift (primary parametrization) | 0.062 | 0.008 | < 10^-13 | 12.3 | 501,092 |
| Gave to ideology alternative (tight band) | -0.161 | 0.017 | < 10^-19 | 16.1 | 747,821 |
| Gave to ideology alternative (median band) | -0.063 | 0.016 | 0.0001 | 14.9 | 699,005 |
| Gave to non-focal same-state winner | 0.011 | 0.022 | 0.62 | 17.6 | 805,812 |

*Note:* Conventional local-linear discontinuities (win side minus loss side) with robust bias-corrected standard errors and confidence intervals, MSE-optimal bandwidths, SEs clustered by focal race. The premium rows are the derived candidate-level outcomes of Section 3; the crossover row is loss-side over win-side (bandwidth selection explains small differences from the Appendix B.4 limit differences). The winnerward-shift row reports the sign-aligned magnitude (Section 3); the composition rows exclude observations lacking party-linked House giving in either cycle (Appendix A.3); negative ideology-alternative estimates mean higher alternative-giving on the losing side.

**Table 4:** Redirection by donor type.

| Donor group | Gave to winner | Gave to loser | \|Partisan reorientation\| | Δ log total giving |
|---|---|---|---|---|
| Individual, small (under $200) | 0.144 (0.014) | -0.089 (0.025) | 0.005 (0.002) | 0.04 (0.15) |
| Individual, mid ($200-500) | 0.300 (0.015) | -0.157 (0.018) | 0.042 (0.004) | 0.33 (0.14) |
| Individual, large (over $500) | 0.472 (0.014) | -0.255 (0.026) | 0.042 (0.004) | 0.48 (0.15) |
| Business-connected PAC | 0.374 (0.020) | -0.213 (0.033) | 0.065 (0.010) | — |
| Labor PAC | 0.367 (0.046) | -0.196 (0.039) | 0.023 (0.016) | — |
| Non-connected PAC (no sponsor) | 0.456 (0.035) | -0.192 (0.038) | 0.039 (0.015) | — |
| Ideological or single-issue (CRP) | 0.451 (0.022) | -0.231 (0.039) | 0.042 (0.008) | — |
| Committees pooled (headline) | 0.375 (0.016) | -0.227 (0.023) | 0.062 (0.008) | 0.03 (0.13) |

*Note:* Race-clustered SEs in parentheses. Individual rows are individual-panel estimates on fixed size strata; committee rows are committee-panel estimates. The three FEC rows are mutually exclusive (sponsor-recovered classification); non-connected = 308 committees with no sponsor. The CRP row classifies purpose directly (bridged subsample; Appendix Section P). Partisan reorientation is the winnerward-shift magnitude, reported without trend adjustment (pre-period trends and trend-adjusted values: Appendix Section J for committees, Appendix Section P for individuals); the labor value is not significant ($p = 0.18$). Δ log total giving is the fourth co-primary outcome; only the mid and large values differ from zero ($p = 0.03, 0.003$).

# References


Akey, P. (2015). Valuing changes in political networks: Evidence from campaign contributions to close congressional elections. *Review of Financial Studies*, 28(11):3188–3223.

Aldrich, J. H., Ballard, A. O., Rohde, D. W., and Lerner, J. (2017). Does the gift keep on giving? house leadership PAC donations before and after majority status. *Journal of Politics*, 79(4):1449–1453.

Ansolabehere, S., de Figueiredo, J. M., and Snyder, J. M. (2003). Why is there so little money in U.S. politics? *Journal of Economic Perspectives*, 17(1):105–130.

Armstrong, T. B. and Kolesár, M. (2018). Optimal inference in a class of regression models. *Econometrica*, 86(2):655–683.

Barber, M. (2016). Donation motivations: Testing theories of access and ideology. *Political Research Quarterly*, 69(1):148–159.

Barber, M. and Eatough, M. (2020). Industry politicization and interest group campaign contribution strategies. *Journal of Politics*, 82(3):1008–1025.

Bonica, A. (2014). Mapping the ideological marketplace. *American Journal of Political Science*, 58(2):367–386.

Bonica, A. (2018). Inferring roll-call scores from campaign contributions using supervised machine learning. *American Journal of Political Science*, 62(4):830–848.

Brunell, T. L. (2005). The relationship between political parties and interest groups: Explaining patterns of PAC contributions to candidates for congress. *Political Research Quarterly*, 58(4):681–688.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Cassie, W. E. and Thompson, J. A. (1998). Patterns of PAC contributions to state legislative candidates. In Thompson, J. A. and Moncrief, G. F., editors, *Campaign Finance in State Legislative Elections*, pages 158–184. CQ Press, Washington, DC.

de la Cuesta, B. and Imai, K. (2016). Misunderstandings about the regression discontinuity design in the study of close elections. *Annual Review of Political Science*, 19:375–396.

De Magalhães, L., Hangartner, D., Hirvonen, S., Meriläinen, J., Ruiz, N. A., and Tukiainen, J. (2025). When can we trust regression discontinuity design estimates from close elections? evidence from experimental benchmarks of the incumbency advantage. *Political Analysis*, 33(3):258–265.

Dumas, N. K. and Shohfi, K. D. (2018). The behavioral consequences of election outcomes: Evidence from campaign contributions. *British Journal of Political Science*. Advance online publication.

Fouirnaies, A. and Hall, A. B. (2014). The financial incumbency advantage: Causes and consequences. *Journal of Politics*, 76(3):711–724.

Fouirnaies, A. and Hall, A. B. (2018). How do interest groups seek access to committees? *American Journal of Political Science*, 62(1):132–147.

Fowler, A., Garro, H., and Spenkuch, J. L. (2020). Quid pro quo? corporate returns to campaign contributions. *Journal of Politics*, 82(3):844–858.

Francia, P. L., Green, J. C., Herrnson, P. S., Powell, L. W., and Wilcox, C. (2003). *The Financiers of Congressional Elections: Investors, Ideologues, and Intimates*. Columbia University Press, New York.

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

Kim, I. S., Stuckatz, J., and Freiheyt, L. W. (2026). Systemic and sequential links between campaign donations and lobbying. *The Journal of Politics*, 88(1):110–124.

Lee, D. S. (2008). Randomized experiments from non-random selection in U.S. house elections. *Journal of Econometrics*, 142(2):675–697.

Lee, D. S., Moretti, E., and Butler, M. J. (2004). Do voters affect or elect policies? Evidence from the U.S. House. *Quarterly Journal of Economics*, 119(3):807–859.

Lewis, J. B., Poole, K., Rosenthal, H., Boche, A., Rudkin, A., and Sonnet, L. (2025). Voteview: Congressional roll-call votes database. `https://voteview.com/`.

Liu, H. (2022). Campaign contributions and access to congressional offices: Patterns in foreign lobbying data. *Political Research Quarterly*, 75(3):812–828.

Marshall, J. (2024). Can close election regression discontinuity designs identify effects of winning politician characteristics? *American Journal of Political Science*, 68(2):494–510.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714.

Meisels, M. (2026). Everything in moderation? The effect of extremist nominations on individual and corporate PAC fundraising. *Political Science Research and Methods*, 14(2):353–369.

Myers, A. C. W. (2026). Do donors punish extremist primary nominees? Evidence from Congress and American state legislatures. *American Political Science Review*, 120(3):1059–1079.

Nokken, T. P. and Poole, K. T. (2004). Congressional party defection in American history. *Legislative Studies Quarterly*, 29(4):545–568.

Phillips, C. H. (2024). Interest group strategy and state legislative polarization. Working paper, Vanderbilt University.

Powell, E. N. and Grimmer, J. (2016). Money in exile: Campaign contributions and committee access. *Journal of Politics*, 78(4):974–988.

Romer, T. and Snyder, Jr., J. M. (1994). An empirical investigation of the dynamics of PAC contributions. *American Journal of Political Science*, 38(3):745–769.

Rueda, M. R. and Ruiz, N. A. (2025). How do electoral outcomes affect campaign contributions? the role of personal loyalty and investment motives. *British Journal of Political Science*, 55:e86.

Snyder, James M., J. (1990). Campaign contributions as investments: The u.s. house of representatives, 1980–1986. *Journal of Political Economy*, 98(6):1195–1227.

Stratmann, T. (2002). Can special interests buy congressional votes? evidence from financial services legislation. *Journal of Law and Economics*, 45(2):345–373.

Stratmann, T. (2005). Some talk: Money in politics. a (partial) review of the literature. *Public Choice*, 124:135–156.

Welch, W. P. (1980). The allocation of political monies: Economic interest groups. *Public Choice*, 35:97–120.

This appendix provides supporting analyses referenced in the body of the manuscript. The numbering of tables and figures matches the cross-references in the body.

# A Data construction details

## A.1 DIME extraction and sample construction

I extract committee contributions to U.S. House general-election candidates from the Database on Ideology, Money in Politics, and Elections (Bonica, 2014, 2018), restricting to `contributor_type = "C"` (committees: the full set of registered committee filers, as distinct from individual donors). The raw DIME file contains every itemized contribution from 1980 onward; I restrict to 1990 forward because committee coverage stabilizes in the early 1990s. The races I study span election cycles 1990 through 2022, with outcomes measured in the following cycle (t+2), through 2024.

DIME's 2024 release is complete on the contribution side, but its recipient records carry no election-result fields (general-election vote shares and winner indicators are empty), which breaks the filter that defines the general-election universe for outcome measurement. Simply widening the 2024 universe to all candidate rows was rejected: 18.8 percent of 2022 committee House dollars went to non-general candidates, so an all-candidate universe would break outcome comparability for exactly the newest cohort. I instead reconstruct 2024 general-election universe membership from the MIT Election Lab House returns, matching on state, district, surname, and party (with squashed-name, accent-stripped, and redistricting-renumbering fallbacks). Validated on 2022, where DIME's native universe is observed, the identical procedure captures 99.65 percent of committee dollars; the residual is dominated by DIME-side classification errors and territorial delegates absent from MIT (roughly 0.05 percent of dollars). Only universe membership is patched (party codes and contribution records remain DIME's own), and outcome cycles through 2022 keep the native universe unchanged. General-election winner flags for the

2024 cycle, required by the geographic-complement outcome, are derived from the same MIT returns and agree with DIME's winner coding in 99.76 percent of 2022 races; both disagreements resolve against DIME on the public record.

For each contested two-candidate U.S. House general election ("the focal race", held in the "focal cycle" t), I identify every committee that contributed to either of the top two finishers in any cycle up to and including the focal cycle. The construction yields 1,806,639 committee-by-race observations across 5,969 contested races and 117,251 unique committees. Two corrections to DIME's recipient file enter the construction. First, sixteen races carry identical vote shares for their top two candidates in DIME's recipient file, which would place every committee in them at a margin of exactly zero; I resolve the winner and the margin from the MIT Election Lab's official vote totals, which recovers nine genuine contests (eight near-ties with official margins between 0.001 and 0.41 percentage points, among them Iowa's second district in 2020, decided by six votes, and one lopsided race carrying a duplicate entry) and drops the seven remaining entries, which are recipient-record artifacts rather than two-candidate contests (a same-party duplicate entry for an unopposed or lopsided seat, or a missing district code). No race enters at a margin of zero. Second, each candidate's incumbency status is taken from external rosters rather than from DIME's candidate-status field: a candidate is a sitting member if Voteview lists the candidate in the House that sat during the focal cycle, cross-checked against candidate continuity in the MIT returns. DIME's own field codes most open-seat contenders in the 2016 through 2020 cycles as challengers and would leave 4, 4, and 9 open-seat races in those cycles, against 38, 59, and 42 under the external rosters (Appendix Section B.2).

## A.2 Pick rule sensitivity

A "dual-pick" PAC×race observation is one where a PAC contributed in cycle t to both top-2 candidates in the same race. Of 14,705 dual-pick PAC×race pairs in the analytic sample (0.8 percent of 1,806,639 PAC×race observations), 3,096 (21.1 percent) involve

equal dollar amounts to both candidates ("tie cases") and 11,609 (78.9 percent) involve unequal amounts. The primary specification applies a dominant-amount rule: keep the row corresponding to the candidate who received more dollars in cycle t, with ties broken lexicographically by `supported`. I assess sensitivity to three alternative rules: (i) drop-dual (exclude all dual-pick PAC×race observations from the analysis); (ii) replicate (include both candidate-rows for each dual-pick PAC×race, with the running variable sign-flipped per candidate); (iii) share-weighted (include both rows, weighted by each candidate's dollar share within the dual-pick PAC×race pair, so tie-amount dual-picks contribute weight 0.5 per row and unequal-amount dual-picks contribute weights proportional to dollar share). The drop-dual sensitivity is implemented inside the 27-specification robustness curve (Section B); the replicate and share-weighted sensitivities are run as separate alternative-pick-rule estimations. Because `rdrobust` does not accept observation weights, the share-weighted estimator is implemented as a weighted local-linear regression at the `rdrobust`-selected MSE-optimal bandwidth with race-clustered standard errors; standard errors under the share-weighted rule are cluster-robust but not Calonico–Cattaneo–Titiunik bias-corrected.

Table A2 reports the four pre-specified co-primary outcomes plus the sign-aligned derivative `delta_toward_winner` under each pick rule. Across the three alternatives and the four pre-specified outcomes, the largest absolute deviation from the primary specification is 1.0 percentage point (drop-dual versus primary, on `gave_to_winner_t2`); the deviation on `gave_to_loser_t2` is at most 0.8 percentage points, and deviations on `delta_dem_share` and `delta_toward_winner` are below 0.3 percentage points. The remaining co-primary outcome, `delta_log_total_giving`, is measured in log units rather than percentage points; its largest deviation is 0.011 log points.

The headline +0.375 winner-redirection effect, the $-0.227$ loser-loyalty drop, the precise null on total giving, the cancellation null on raw `delta_dem_share`, and the

Table A2: Pick-rule sensitivity (5 outcomes $\times$ 4 rules).

Reported point estimates with standard errors in parentheses. The primary, drop-dual, and replicate columns use `rdrobust` (Calonico et al., 2014) with a triangular kernel, MSE-optimal bandwidth, and race-clustered HC1 standard errors. The share-weighted column uses a weighted local-linear regression at the `rdrobust`-selected MSE-optimal bandwidth with race-clustered standard errors via `fixest::feols`. The `delta_toward_winner` row reports the absolute value $|\beta|$, matching the manuscript-primary parametrization (Section 3 footnote; Appendix A.3).

| Outcome | Primary (dominant-amount) | Drop-dual | Replicate | Share-weighted |
|---|---|---|---|---|
| `gave_to_winner_t2` | +0.375 (0.016) | +0.385 (0.015) | +0.371 (0.016) | +0.377 (0.012) |
| `gave_to_loser_t2` | −0.227 (0.023) | −0.229 (0.023) | −0.219 (0.023) | −0.225 (0.020) |
| `delta_log_total_giving` | +0.032 (0.131) | +0.043 (0.131) | +0.040 (0.131) | +0.039 (0.084) |
| `delta_dem_share` | +0.010 (0.009) | +0.011 (0.009) | +0.010 (0.009) | +0.011 (0.008) |
| `delta_toward_winner` ($\lvert\beta\rvert$) | +0.062 (0.008) | +0.062 (0.008) | +0.060 (0.008) | +0.062 (0.006) |

+6.2 percentage-point sign-aligned `delta_toward_winner` discontinuity are jointly preserved across all four pick rules. The 27-specification robustness curve (Section B) provides additional pick-rule-independent coverage.

## A.3 Outcome construction

The four co-primary outcomes are constructed at cycle t+2. `gave_to_winner_t2` is an indicator for any positive contribution from the PAC to the focal-race winner, and `gave_to_loser_t2` the same indicator for the focal-race loser. `delta_log_total_giving` is log(1 + total giving at t+2) minus log(1 + total giving at t). `delta_dem_share` is the Democratic share of the PAC's House giving at t+2 minus the same share at t, the pre-specified raw partisan-composition outcome, signed in the Democratic direction.

I additionally report the sign-aligned derivative `delta_toward_winner` defined as `delta_dem_share` multiplied by an indicator that takes value +1 in races won by a Democrat and -1 in races won by a Republican. The two parametrizations are a one-to-one transform of each other at the observation level and add no information, but the sign alignment does change the pooled test: it removes the opposite-sign cancellation across Democratic-won and Republican-won pools, which is why the raw outcome

shows no pooled discontinuity (+0.010, standard error 0.009) while the aligned one does (Table 3). The sign-aligned derivative gives the partisan-composition outcome the sign convention of the framework's PAC-perspective treatment indicator (positive value = movement toward the focal-race winner's party in every pool) and is the manuscript's primary parametrization for this outcome. The 27-specification robustness curve already uses the same sign-normalization convention for `logit_rdd` and `race_margin_lpm`. Family-wise correction (Bonferroni and BH) is applied to the four pre-specified outcomes; `delta_toward_winner` is reported alongside as a deterministic transform of `delta_dem_share`, not as a fifth test. Both composition outcomes are defined only when the committee has Democratic or Republican House giving in both cycles: 248,158 committee-race observations (13.7 percent) have none at t+2 (overwhelmingly committees inactive at t+2), a further 20 have none at t, and a further 1,006 are in races won by a third-party candidate, so `delta_toward_winner` is missing on 249,184 rows (13.8 percent). This is why the composition rows of Table 3 have smaller effective samples than the total-giving row, and it makes the composition outcomes conditional on remaining active at t+2; the discontinuity in t+2 activity at the cutoff is 1.1 percentage points (Appendix Section J), which bounds the selection this conditioning can introduce.

The inverse-hyperbolic-sine transformation of total giving yields a discontinuity of +0.044 (SE 0.138; Table Q1, row B), confirming the log-handling does not drive the null.

## A.4 Assignment structure and effective sample size

The running variable is a race-level quantity: every committee-by-race observation attached to one focal race carries the same win-loss margin, signed per committee by the side it backed, so treatment is assigned once per race and the 1,806,639 committee-by-race observations inherit 5,969 independent close-race draws. Identification is unaffected by this repetition. Continuity of potential outcomes in the margin at the cutoff is a race-level condition, and each committee inherits its race's quasi-random assignment. The

Table A3: Distinct focal races within each outcome's MSE-optimal bandwidth.

| Outcome (Table 3 row) | h (pp) | Races within bandwidth |
|---|---|---|
| Gave to focal winner | 15.1 | 1,682 |
| Gave to focal loser | 6.6 | 743 |
| Funded own pick (continuation premium) | 6.5 | 737 |
| Funded non-picked candidate (crossover premium) | 9.0 | 966 |
| Change in log total House giving | 15.7 | 1,744 |
| Change in Democratic share (pre-specified) | 11.4 | 1,260 |
| Winnerward shift (primary parametrization) | 12.3 | 1,355 |
| Gave to ideology alternative (tight band) | 16.1 | 1,784 |
| Gave to ideology alternative (median band) | 14.9 | 1,635 |
| Gave to non-focal same-state winner | 17.6 | 1,958 |

*Note:* Distinct focal races with $|\text{rv}| \leq \text{h}$ in the analytic panel, at each outcome's Table 3 bandwidth. Complete-case filtering within the window changes these counts by at most eight races (largest: winnerward shift, 1,355 to 1,347), because a race drops only when every committee row attached to it is missing the outcome. The full panel spans 5,969 focal races.

repetition matters for inference: observations within a race share the assignment and its shocks, so all standard errors in the paper cluster at the focal race, which Section 3 reports is roughly three times more conservative than clustering at the PAC. Table A3 counts the distinct focal races inside each outcome's MSE-optimal bandwidth; every estimate in Table 3 rests on at least 737 race clusters. Because margins repeat across the committees in a race, the running variable has mass points; bandwidth selection and variance estimation use the estimator's default mass-points adjustment (`masspoints = adjust`, `rdrobust` 3.0.0).

The committee-by-race unit also fixes the estimand's weighting. Effects average over donor–candidate relationships, so a race weighs in proportion to the committees that gave in it; a race-weighted estimand would answer a different question (the effect on the average race rather than on the average giving relationship), and the paper's outcomes are properties of those relationships. The same structure underlies the observation-level density asymmetry of Appendix Table J1: committees that backed the eventual winner are mechanically more numerous, so the density test rejects for compositional reasons rather than manipulation, and validity rests on the race-clustered covariate balance of Section 3 (Appendix Section J).

# B Specification curve

## B.1 Estimator-perturbation grid

Table B1 reports the full 27-specification grid for the headline `gave_to_winner_t2` outcome. Its era buckets are assigned on the focal cycle. The first two labels name their focal cycles directly; the third, `2016–2024`, runs from its first focal cycle to its last outcome cycle, covering focal cycles 2016 through 2022, the last in the panel. These are the labels the `era` variable carries in the analysis panel, so the table and the replication artifact agree.

Table B1: Estimator-perturbation robustness grid (27 specifications).

| Specification | Estimate | SE | p-value | h (pp) |
|---|---|---|---|---|
| primary | 0.375 | 0.016 | < 10^-100 | 15.1 |
| h = 1 | 0.391 | 0.049 | 1.3e-15 | 1 |
| h = 2 | 0.343 | 0.043 | 1.2e-18 | 2 |
| h = 3 | 0.347 | 0.038 | 5.3e-20 | 3 |
| h = 4 | 0.357 | 0.034 | 1.4e-22 | 4 |
| h = 5 | 0.363 | 0.032 | 1.4e-26 | 5 |
| h = 7 | 0.367 | 0.029 | 6.6e-35 | 7 |
| h = 10 | 0.374 | 0.025 | 2.2e-48 | 10 |
| h = 15 | 0.375 | 0.020 | 2.2e-75 | 15 |
| h = 20 | 0.374 | 0.017 | 3.7e-103 | 20 |
| CER-optimal | 0.373 | 0.019 | 3.7e-90 | 9.5 |
| Polynomial p = 2 | 0.380 | 0.019 | 2.4e-91 | 22.7 |
| Donut 0.25 | 0.372 | 0.016 | 6.4e-119 | 15.5 |
| Donut 0.50 | 0.376 | 0.018 | 1.2e-102 | 14.1 |
| Donut 1.00 | 0.403 | 0.019 | 2.7e-100 | 10.4 |

| Specification | Estimate | SE | p-value | h (pp) |
|---|---|---|---|---|
| Donut 1.50 | 0.396 | 0.023 | 6.3e-69 | 10.3 |
| Donut 2.00 | 0.380 | 0.024 | 6.6e-58 | 12.0 |
| Cluster = PAC | 0.369 | 0.005 | 0.000 | 7.8 |
| Covariate-adjusted | 0.372 | 0.015 | 2.8e-134 | 15.8 |
| Pick: drop dual | 0.385 | 0.015 | 2.1e-150 | 15.7 |
| Logit RDD (log-odds) | -2.05 | 0.075 | 8.7e-166 | 15.1 |
| Honest M = 0.1 | 0.390 | 0.007 | 0.000 | — |
| Honest M = 0.5 | 0.392 | 0.008 | 0.000 | — |
| Era 1990–2002 | 0.463 | 0.027 | 1.9e-68 | 13.0 |
| Era 2004–2014 | 0.316 | 0.021 | 6.0e-51 | 14.9 |
| Era 2016–2024 | 0.377 | 0.026 | 4.0e-47 | 21.7 |
| Race-margin LPM | -0.374 | 0.011 | 8.3e-237 | 15.1 |

All 27 specifications are sign-consistent (after sign-normalizing the logit and race-margin LPM rows) and significant at the five percent level; the estimates range from 0.32 to 0.46, with a coefficient of variation of 0.07. Only the pre-BCRA era subsample produces a confidence interval that essentially does not overlap the pooled-primary CI.

## B.2 Decompositions by PAC size, PAC ideology, incumbency type, and industry

Per-subgroup rdrobust on `gave_to_winner_t2`.

**PAC size (terciles of lifetime House giving)**

| Tercile | Estimate |
|---|---|
| Small | 0.267 |
| Medium | 0.400 |
| Large | 0.481 |

Monotone in size. All three terciles are strongly positive, ruling out mega-PAC dominance.

**PAC ideology (terciles of the committee's DIME CFscore, `contributor.cfscore`)**

| Tercile | Estimate |
|---|---|
| D-leaning | 0.422 |
| Mixed | 0.316 |
| R-leaning | 0.369 |

Strongly positive across the ideological spectrum; the "mixed" group is mildly attenuated but stays well above zero.

**Incumbency type at focal race**

| Type | Estimate (SE) | Era-adjusted | Focal 1990–2014 |
|---|---|---|---|
| Open seat | 0.435 (0.023) | 0.429 (0.022) | 0.442 (0.024) |
| Incumbent re-elected | 0.442 (0.017) | 0.436 (0.016) | 0.450 (0.019) |
| Challenger wins | 0.373 (0.029) | 0.380 (0.032) | 0.380 (0.033) |

The open-seat and incumbent-reelected estimates are the same to within sampling error: open seats +0.435 against +0.442 for re-elected incumbents, a difference of $-0.007$ with a standard error of 0.028 (era-adjusted $-0.007$, standard error 0.027; the disjoint subsamples are treated as independent), where an information or viability-updating account, which requires prior incumbency to update on, predicts a smaller open-seat effect. Incumbency is classified from external rosters (Appendix Section A.1): DIME's own candidate-status field codes most open-seat contenders in focal 2016–2020 as challengers and would leave 4, 4, and 9 open-seat races in those cycles against 38, 59, and 42 under the

external classification. The 49 races in which both candidates were sitting members (post-redistricting pairings) form a fourth category too small to estimate separately and enter none of the three rows. Under the external classification the open-seat subsample is not concentrated in the earliest era (38 percent of its kernel weight falls in 1990–2002, against 26 percent for incumbent-reelected races), and the era-adjusted column (rdrobust with era indicators as covariates) moves each estimate by less than 0.01. The ordering is not uniform across eras: within 1990–2002 the incumbent-reelected estimate is the largest of the three (+0.520 against +0.473 for open seats), within 2004–2014 open seats lead (+0.405 against +0.394), and within 2016–2024 re-elected incumbents lead again (+0.428 against +0.410).

**Industry (regulated / non-regulated)**

| Industry | Estimate |
|---|---|
| Utilities | 0.336 |
| Financial services | 0.324 |
| Telecommunications | 0.276 |
| Defense | 0.357 |
| Non-regulated | 0.381 |

Every regulated industry returns a strongly positive estimate, and the non-regulated baseline is comparable; the access channel is broad-based across the committee universe rather than concentrated in any one regulated sector.

## B.3 Substantive design alternatives

Where Table B1 perturbs a single estimator, Table B3 varies the design decisions that define the study: how the outcome is operationalized, which subsample is analyzed, how the panel is constructed, how each committee's pick is identified, and the incumbency configuration of the pick. Each row changes one design choice from the headline specification; all use the MSE-optimal bandwidth, a triangular kernel, and race-clustered stan-

dard errors unless the row states otherwise. The pick-rule rows re-stream the raw focal-cycle contributions to reconstruct the pre-deduplication committee-race table (with both rows present for committees that gave to both focal candidates) and re-resolve the pick under each rule; the race-fixed outcome is reused.

Table B3: Substantive design alternatives.

| Dimension | Specification | Estimate | SE | p-value | h (pp) |
|---|---|---|---|---|---|
| Outcome | Extensive margin: gave to winner (= primary) | 0.375 | 0.016 | < 10^-100 | 15.1 |
| Outcome | Log dollars given to winner | 2.995 | 0.127 | 1.3e-126 | 15.3 |
| Outcome | Winner's share of t+2 House giving (active committees) | 0.033 | 0.002 | 3.5e-67 | 23.3 |
| Subsample | Incumbent re-elected | 0.442 | 0.017 | 4.1e-154 | 22.7 |
| Subsample | Open seat | 0.435 | 0.023 | 1.2e-80 | 19.9 |

| Dimension | Specification | Estimate | SE | p-value | h (pp) |
|---|---|---|---|---|---|
| Subsample | Challenger won | 0.373 | 0.029 | 8.1e-39 | 10.5 |
| Subsample | Drop post-redistricting cycles (1992, 2002, 2012, 2022) | 0.369 | 0.019 | 3.0e-88 | 14.6 |
| Subsample | Committees in two or more focal races that cycle | 0.386 | 0.016 | 2.6e-130 | 15.2 |
| Subsample | Committees in a single focal race that cycle | 0.189 | 0.012 | 1.4e-54 | 24.3 |
| Subsample | Drop party and candidate commit-tees | 0.377 | 0.018 | 5.8e-103 | 13.8 |

| Dimension | Specification | Estimate | SE | p-value | h (pp) |
|---|---|---|---|---|---|
| Panel | Race-level collapse (one obs per race-side, committee-count weighted) | 0.375 | 0.016 | 2.0e-131 | 15.1 |
| Panel | Contribution-weighted (by focal-cycle dollars to the pick) | 0.473 | 0.022 | 2.0e-102 | 18.3 |
| Panel | Standard errors clustered on state | 0.375 | 0.018 | 6.0e-99 | 14.7 |
| Panel | Standard errors clustered on cycle | 0.375 | 0.026 | 4.0e-49 | 17.8 |
| Pick rule | Dominant amount (= primary) | 0.379 | 0.016 | 6.4e-131 | 15.1 |

| Dimension | Specification | Estimate | SE | p-value | h (pp) |
|---|---|---|---|---|---|
| Pick rule | Transaction-count plurality | 0.377 | 0.016 | 6.3e-127 | 14.9 |
| Pick rule | First dollar (earliest contribution) | 0.371 | 0.016 | 1.7e-117 | 14.7 |
| Pick incumbency | Pick was the sitting incumbent | 0.330 | 0.023 | 1.6e-49 | 15.3 |
| Pick incumbency | Pick opposed a sitting incumbent | 0.475 | 0.024 | 5.6e-90 | 12.2 |
| Pick incumbency | No incumbent among the race's top two | 0.435 | 0.023 | 1.2e-80 | 19.9 |

Every specification on the binary outcome is positive and significant at the one percent level or better, spanning the range 0.189 to 0.475, and the dollar-denominated and compositional outcomes are likewise strongly positive. The one binary estimate well below the headline, single-race committees at +0.189, remains economically large and highly

significant; the contribution-weighted panel, which weights each committee-race by its cycle-t dollars, is +0.473 (standard error 0.022), so the dollar-weighted redirection is if anything larger than the row-weighted one. The single-race attenuation reflects the mechanical narrowing of a less active committee's giving portfolio: a committee giving to fewer candidates is less likely to give to any specific one, the winner included, which is the same breadth channel the compositional share outcome nets out, and it also tracks the strategic-donor gradient of Section 4.3. The drop-party-and-candidate row retains committees whose FEC type is unclassified; the PAC-only row of Appendix Table P2 (+0.357) restricts to the FEC-classified subsample, which is why the two rows differ. The three pick rules agree to within 0.008 because the transaction-count and first-dollar rules reassign only 1,581 and 6,781 of 1,806,639 committee-races respectively (fewer than one in two hundred), since only committees that split their giving between both focal candidates can be assigned differently.

The pick-incumbency rows split the panel by the incumbency configuration of the close contest: the pick is a sitting incumbent, the pick's top-two opponent is a sitting incumbent, or neither top-two finisher is one. Unlike the incumbent-reelected and challenger-won subsample rows, which condition on the realized winner's type, this classification is pre-determined at the time of the contribution. Sitting status comes from the external rosters of Appendix Section A.1 (Voteview membership in the House that sat during the focal cycle, cross-checked against candidate continuity in the MIT returns), applied to both candidates; where both are sitting members (49 post-redistricting pairings) the pick's status governs. DIME's own per-candidate codes, with a twelve-race hand-audit override for open seats whose winner DIME had miscoded as incumbent, give the same three subsamples as a sensitivity (+0.32, +0.49, and +0.45 against +0.33, +0.47, and +0.43 under the external rosters). The two codings differ on 1.7 percent of committee-race rows, and the one systematic disagreement is that the external rosters count three winners of special elections held on the same day as the general election, and

seated afterwards, as sitting members (875 rows in three races), where the hand audit had treated those seats as vacant. The 2022 cycle is classified by the same external rule. Committees whose pick was the sitting incumbent (the configuration with the deepest pre-existing relationship to protect) still redirect at +0.33 when that incumbent narrowly loses, while committees that backed a non-incumbent redirect at +0.43 to +0.47; the redirection is not an artifact of any one pick configuration.

## B.4 Candidate-level office premiums

The two derived outcomes of Section 3 recombine the pre-specified winner and loser outcomes by candidate rather than by office: on the win side of the cutoff the pick outcome coincides with the winner outcome observation-by-observation, on the loss side with the loser outcome, and conversely for the opponent outcome (the coincidence is exact, verified at a fixed common bandwidth, where the side-specific limits of the derived and source outcomes agree to four decimal places). Table B4 reports their discontinuities (the continuation and crossover premiums of Section 4.1) pooled and split by the contest-configuration coding of Appendix B.3: the pick is a sitting incumbent (I), the pick challenges a sitting incumbent (C), or no incumbent stands (O). Re-running the estimator-perturbation grid of Appendix B.1 on the two premiums (twenty-seven specifications) leaves both signs intact in every cell: the continuation premium runs from +0.09 to +0.33 and falls below +0.15 only where the bandwidth is three points or less (fixed bandwidths of one to three points and PAC-level clustering, whose MSE bandwidth is 2.0 points) or under the two RDHonest smoothness bounds (+0.13 and +0.09); the crossover premium runs from +0.012 to +0.11 and is not distinguishable from zero only at fixed bandwidths of four points or less. Donut, quadratic, covariate-adjusted, dual-pick-dropped, era, logit, and race-margin specifications all reproduce the Table 3 signs.

Table B4: Candidate-level office premiums, pooled and by contest configuration.

| Sample | Continuation premium | SE | p-value | h (pp) | Crossover premium | SE | p-value | h (pp) |
|---|---|---|---|---|---|---|---|---|
| Pooled | 0.200 | 0.031 | < 10^-8 | 6.5 | 0.069 | 0.010 | < 10^-9 | 9.0 |
| Pick is the incumbent (I) | 0.212 | 0.042 | < 10^-5 | 7.1 | 0.096 | 0.015 | < 10^-8 | 10.7 |
| Pick challenges an incumbent (C) | 0.211 | 0.041 | < 10^-5 | 9.4 | 0.012 | 0.016 | 0.55 | 10.1 |
| Open seat (O) | 0.295 | 0.048 | < 10^-8 | 10.7 | 0.084 | 0.012 | < 10^-11 | 16.0 |

*Note:* rdrobust conventional local-linear estimates with robust bias-corrected standard errors and confidence intervals, MSE-optimal bandwidths per outcome and sample, triangular kernel, race-clustered SEs. Crossover premiums are reported loss-side over win-side, so positive values mean a candidate's victory attracts money from the committees that opposed them. The side limits behind the four levels quoted in Section 4.1 are computed at each source outcome's canonical bandwidth (15.1 pp for the winner outcome, 6.6 pp for the loser outcome): 0.477 (win side, own pick), 0.028 (win side, opponent), 0.255 (loss side, own pick), and 0.102 (loss side, opponent); their differences (0.222 and 0.074) differ from the premium estimates only through per-outcome bandwidth selection. At the premium regressions' own MSE-optimal bandwidths (6.5 and 9.0 pp), the side limits are 0.456 (win side, own pick), 0.256 (loss side, own pick), 0.026 (win side, opponent), and 0.095 (loss side, opponent); their differences reproduce the Table 3 premiums (0.200, 0.069) exactly. Within-configuration estimates are causal for that subpopulation; comparisons across configurations are descriptive, because committees select which candidates to back and the configurations therefore hold different donor populations.

## B.5 Winner position stability at the margin

The crossover premise of Section 4.1 has two parts: the opponent's proximity to the committee is the same on both sides of the cutoff before the vote, and narrow winners do not moderate once in office (Lee et al., 2004). The second part cannot be tested for the opponent directly, because roll-call positions exist only for winners and the committee-side CFscore at $t + 2$ is estimated from the $t + 2$ giving that is the outcome. It can be checked for the winners of the focal races. I bridge each winner's DIME identifier to its ICPSR number through the DIME recipient file and attach the Congress-specific (Nokken–Poole) first-dimension score from Voteview (Nokken and Poole, 2004; Lewis et al., 2025); 5,754

of the 5,969 focal-race winners bridge to a House member who holds a score at the post-focal Congress and whose state matches the race (every bridged member state-matches). Movement is signed so that positive values denote movement toward the opponent's party.

Winners with a score in both the pre-focal and the post-focal Congress (4,628 re-elected members facing a candidate of the other party) do not move. Their change in score, signed toward the opponent, averages $+0.006$ (SE 0.005) at margins under 5 points, $-0.001$ (SE 0.002) under 15 points, and $-0.001$ (SE 0.001) at all margins, with a slope on margin of $-0.0003$ per point (SE 0.0005) within 15 points; the mean absolute change is 0.053 at margins under 5 points against 0.049 for all 6,278 continuing House members across the same Congress pairs (SEs clustered by member). Freshman winners have no pre-focal score; their first-term score relative to the party–Congress median, signed toward the center, is $+0.084$ (SD 0.140) at margins under 5 points, against $+0.088$ (SD 0.137) for re-elected winners at the same margins (difference $-0.004$, SE 0.013; $-0.003$, SE 0.007, within 15 points). The moderation of narrow winners is a property of the districts that produce narrow races, shared by freshmen and incumbents alike, not a first-term shift.

Table B5 asks whether the small movements that do occur are followed by money. Among committees whose pick lost to a re-elected winner (an incumbent), the loss-side crossover rate at $t + 2$ is flat across terciles of the winner's movement toward the committee's side (side is the sign of the committee CFscore minus the winner CFscore, both on the DIME scale; tercile means $-0.076$, $-0.003$, and $+0.070$): 3.4, 4.3, and 3.3 percent at margins under 5 points (T3 minus T1: $-0.1$ points, SE 0.9) and 3.0, 3.6, and 2.8 percent under 10 points. On the win side, committees whose incumbent pick won continue funding at 46.0, 52.9, and 50.2 percent at margins under 5 points (T3 minus T1: $+4.2$ points, SE 2.1), $+2.4$ points (SE 1.4) under 10 points, and $+0.3$ (SE 0.6) at all margins; the gradient is not monotone and shrinks as the bandwidth widens. Both panels condition on a post-treatment trait of the winner and are descriptive (Marshall, 2024; the PAC-side mod-

Table B5: Giving at $t + 2$ by tercile of the winner's in-office movement toward the committee.

| Side and outcome | Margin | T1 (away) | T2 | T3 (toward) | T3 − T1 (SE) | Obs. | Races |
|---|---|---|---|---|---|---|---|
| Loss side: crossover | < 5 pp | 3.4 | 4.3 | 3.3 | $-0.1$ (0.9) | 26,030 | 233 |
| | < 10 pp | 3.0 | 3.6 | 2.8 | $-0.2$ (0.5) | 53,387 | 530 |
| | all | 3.1 | 3.4 | 2.7 | $-0.5$ (0.4) | 124,092 | 4,149 |
| Win side: continuation | < 5 pp | 46.0 | 52.9 | 50.2 | $+4.2$ (2.1) | 89,615 | 239 |
| | < 10 pp | 47.3 | 49.1 | 49.7 | $+2.4$ (1.4) | 192,527 | 537 |
| | all | 50.7 | 52.4 | 50.9 | $+0.3$ (0.6) | 1,270,885 | 4,737 |

*Note:* Committee-by-race rows whose race winner holds Nokken–Poole scores in both the pre-focal and the post-focal Congress (a re-elected incumbent) and whose committee and winner CFscores are both observed (1,394,977 rows, 4,737 races). Movement is the winner's post-focal score minus the pre-focal score, multiplied by the sign of the committee CFscore minus the winner CFscore, so that positive values denote movement toward the committee's side; terciles are formed within each side (loss-side breaks $-0.028$ and $+0.022$, win-side breaks $-0.025$ and $+0.023$). Cells are the percentage of committees giving at $t + 2$ to the winning opponent (loss side) or to the winning pick (win side) within the margin window; T3 − T1 is in percentage points with a race-clustered SE. These are rates within a side, not discontinuities: because the movement is defined only when the incumbent won, the win side holds committees that backed the incumbent (configuration I of Table B4) and the loss side committees that backed the challenger (configuration C), and a discontinuity across the two would compare different committee populations. Script: `code/65_position_stability.R`.

erator of Section 4.4 has the same status), and no within-tercile discontinuity is defined: the movement exists only when the incumbent won, so on the win side the pick is the incumbent and on the loss side the pick is the challenger. The crossover channel, the one the premise concerns, does not respond to the winner's in-office movement; the continuation gradient at the narrowest margins is at most a small ideological increment on a base near one half.

# C Bandwidth sensitivity

Estimates across the bandwidth grid h in {1, 2, 3, 4, 5, 7, 10, 15, 20} appear in Table B1 above. The bandwidth grid produces a stable estimate range [0.34, 0.39] for the primary outcome, with no monotone trend in either direction. The remaining co-primary outcomes (`gave_to_loser_t2`, `delta_log_total_giving`, and `delta_dem_share`) show similarly stable patterns across the same grid.

# D Donut-hole sensitivity

The donut-hole specification drops observations within $|rv| < d$ for d in {0.25, 0.50, 1.00, 1.50, 2.00} percentage points. Estimates range from 0.372 to 0.403, each at least fifteen standard errors from zero (Table B1). The donut-hole grid addresses the precise-manipulation alternative (in which manipulation is concentrated within a narrow window of the cutoff); the absence of a monotone or large pattern across donut sizes argues against precise manipulation as a threat.

# E Era subsamples and Wald interaction

Table E1 re-estimates the era split at bandwidths selected per cell for the balanced-PAC comparison; the aggregate era estimates the body reports (Section 4.5) are the Appendix Table B1 rows, which coincide with E1's full-sample column now that both use the same race-clustered HC1 variance estimator (bandwidths are selected per cell). The access redirection weakens over time but remains substantial in every era, and the formal era $\times$ loss interaction is well outside chance. Table E1 reports the era split for both the full sample and a balanced-PAC subset restricted to the 1,262 committees (1.1 percent of the 117,251 unique committees in the panel) observed in all three eras. The balanced-PAC comparison addresses the adversarial-reviewer concern that the era heterogeneity could partly reflect compositional shifts in DIME's committee coverage rather than a behavioral change in PAC redirection (the coverage drift is documented in Bonica, 2014).

Significance: *p<0.05, **p<0.01, ***p<0.001 (two-sided, conventional rdrobust inference for the row-level coefficients; t-statistics for the $\Delta$ row are derived against the full-sample SE under independence, conservative because the balanced subset is a strict subsample).

Two findings emerge from the side-by-side comparison. First, the *era decline magnitude* tracks closely across the two samples. From 1990–2002 to 2004–2014, the full-sample

Table E1: Era subsamples on `gave_to_winner_t2`, full sample vs balanced-PAC subset.

Coefficients are rdrobust point estimates with race-clustered HC1 standard errors in parentheses below; bandwidth selected per cell by MSE-optimal CCT mserd; triangular kernel; local-linear polynomial. The balanced subset restricts to PACs with $\geq$1 PAC $\times$ focal-race observation in each of the three era buckets (`era ∈ {1990–2002, 2004–2014, 2016–2024}`). $\Delta$ rows report the level shift between the balanced subset and the full sample, with the in-SE distance computed against the full-sample standard error.

| | 1990–2002 (pre-BCRA) | 2004–2014 | 2016–2024 (post-Citizens-United) |
|---|---|---|---|
| **Full sample** | | | |
| $\hat{\beta}_{\text{loss}}$ | 0.463*** | 0.316*** | 0.377*** |
| | (0.024) | (0.019) | (0.021) |
| Bandwidth h (pp) | 13.0 | 14.9 | 21.7 |
| N obs (within h) | 167,086 | 270,516 | 320,068 |
| N PACs (within h) | 8,804 | 35,930 | 38,052 |
| **Balanced-PAC subset** | | | |
| $\hat{\beta}_{\text{loss}}$ | 0.532*** | 0.385*** | 0.384*** |
| | (0.027) | (0.028) | (0.043) |
| Bandwidth h (pp) | 14.6 | 12.3 | 13.4 |
| N obs (within h) | 35,215 | 49,147 | 33,769 |
| N PACs (within h) | 962 | 1,164 | 967 |
| $\Delta$ **(balanced** $-$ **full)** | +0.069 | +0.069 | +0.008 |
| (full-sample SE units) | (+2.9 SE)** | (+3.7 SE)*** | (+0.4 SE) |

redirection drops by 0.147 on the binary outcome (0.463 to 0.316); over the same span the balanced-PAC subset drops by 0.147 (0.532 to 0.385). From 1990–2002 to 2016–2024, the full-sample decline is 0.087 (0.463 to 0.377) and the balanced-PAC decline is 0.148 (0.532 to 0.384). The recent-era recovery is weaker among long-lived committees, so the balanced subset declines somewhat more over the full span. The formal era $\times$ loss interaction rejects a common effect across eras (Wald F statistic 13.94), so the cross-era difference is not chance. The first-transition gradient, the pre-BCRA-to-post-BCRA drop on which the era argument depends, survives essentially unchanged when DIME's coverage drift is held constant by restricting to PACs continuously observed across all three eras, which attributes the era heterogeneity to behavioral change rather than to compositional turnover in the coverage panel.

Second, the *level* of redirection is higher in the balanced-PAC subset in the first two eras, by seven percentage points (2.9 and 3.7 full-sample standard errors respectively, both outside conventional confidence intervals), while in 2016–2024 the two samples are

statistically indistinguishable (+0.8 percentage points, 0.4 standard errors). Long-lived committees were more access-driven than the cross-section average through 2014; the panel's average effect in those eras is in part diluted by the entry and exit of more transient PACs whose redirection is somewhat smaller, and by the latest era the gap has closed. The level pattern is a substantive finding about committee heterogeneity rather than a threat to the era decline, whose first transition operates within these long-lived committees at essentially the magnitude observed in the panel.

The sign-aligned outcome `delta_toward_winner` (rdrobust returns the right limit minus the left limit, so the effect of a loss is the sign-flipped value) shows almost no full-vs-balanced separation: era estimates are $-0.054$ / $-0.082$ / $-0.047$ in the full sample and $-0.053$ / $-0.080$ / $-0.042$ in the balanced subset. The partisan-composition channel is essentially flat across both era and sample.

## E.1 Within-vs-cross-party decomposition

Two readings of the era attenuation are a priori plausible (body Section 4.5): post-2002 regulatory channel substitution (the BCRA soft-money ban and post-Citizens-United 501(c)(4) growth absorbing access-seeking dollars into vehicles not captured by direct PAC giving) and polarization-driven dampening of cross-aisle redirection. The two readings make opposite predictions about within-vs-cross-party PACs: a polarization story predicts the cross-party component (PAC ideology mismatched with the winner's party) to shrink most, since polarization makes cross-aisle support less palatable; a channel-substitution story predicts no particular cross-vs-within asymmetry on its own, but is consistent with within-party PACs migrating earlier into alternative vehicles because their access alignment is more substitutable across channels.

I classify each PAC $\times$ focal-race observation as `within_party` if the PAC's CF-score sign matches the focal-race winner's party affiliation (D-leaning PACs paired with D winners; R-leaning PACs paired with R winners) and `cross_party` if mis-

Table E2: Within-vs-cross-party era split on `gave_to_winner_t2`.

rdrobust point estimate (SE), MSE-optimal bandwidth selected per cell, race-clustered HC1 SE, triangular kernel.

| Era | Within-party | Cross-party |
|---|---|---|
| 1990–2002 (pre-BCRA) | +0.401 (0.049) | +0.483 (0.041) |
| 2004–2014 | +0.174 (0.044) | +0.430 (0.034) |
| 2016–2024 (post-Citizens-United) | +0.267 (0.064) | +0.428 (0.052) |

matched. The sample is restricted to non-null PAC CFscore (98.6 percent of the panel; DIME's `contributor.cfscore` is the PAC-ideology measure throughout this paper) and to focal races where `winner_party` $\in$ {`100=Dem, 200=Rep`}. The decomposition is run on the four pre-specified co-primary outcomes plus the sign-aligned `delta_toward_winner` (30 cells total: 5 outcomes $\times$ 3 eras $\times$ 2 components).

The within-party trajectory (+0.401, +0.174, +0.267) drops 57 percent from the pre-BCRA era to 2004–2014, with partial recovery in 2016–2024. The cross-party trajectory (+0.483, +0.430, +0.428) falls by 0.053 over the first transition, about one cell standard error, and is flat thereafter. The era attenuation in the pooled `gave_to_winner_t2` estimator is therefore mainly a within-party phenomenon: the within-party drop from 1990–2002 to 2004–2014 (0.228) is roughly four times the cross-party drop over the same span (0.053).

This pattern disposes of one polarization story but not another. It rules out the version in which rising polarization makes cross-aisle support less palatable, which predicts the cross-party component to shrink most, whereas the cross-party component is the stable one. But a different polarization mechanism (polarization tightening within-party loyalty, so that committees increasingly default to their own copartisans regardless of the focal race) predicts exactly the observed pattern (a shrinking within-party redirection with a stable cross-party one) and is observationally equivalent here to channel substitution, in which within-party PACs already aligned with the focal-race winner's party migrated marginal access flows into alternative vehicles after the BCRA soft-money ban

while cross-party PACs continued direct-PAC redirection because the cross-aisle access signal is harder to deliver through 501(c)(4) and super-PAC channels (which are typically aligned with the PAC's home-party network rather than the rival party's). The within-cross decomposition therefore rules out the cross-aisle-palatability version of polarization but leaves channel substitution and within-party-loyalty tightening as the two live and unseparated readings; cleanly separating them at the mechanism level would require state-level variation in campaign-finance rules within the post-BCRA window, which I leave to future work (body Section 4.5).

The full decomposition data additionally reports analogous estimates for the remaining co-primary outcomes (`gave_to_loser_t2`, `delta_log_total_giving`, `delta_dem_share`) and the sign-aligned derivative `delta_toward_winner`. Notable patterns from the secondary outcomes: (i) within-party loser-loyalty drop is consistently larger (more negative) than cross-party loser-loyalty drop in every era (1990–2002: $-0.315$ vs $-0.221$; 2004–2014: $-0.189$ vs $-0.128$; 2016–2024: $-0.288$ vs $-0.145$), substantively coherent with within-party PACs having a stronger access incentive to follow the winner of "their" party; (ii) `delta_dem_share` shows opposite signs across components by construction (within-party PACs by definition align with the winner's party direction); (iii) `delta_toward_winner` shows similar magnitude winnerward shifts across components within each era (no within-cross asymmetry on the sign-normalized partisan-composition outcome), consistent with both kinds of PACs moving toward the focal-race winner's party in the loser-side pool relative to the winner-side pool.

## E.2 Cross-party decomposition by sector, geography, and ideological distance

The within-vs-cross-party split of E.1 shows that the cross-party component is the era-stable, model-diagnostic half of the redirection. This subsection pools that component across the three eras and decomposes it along three predetermined moderators, sharpen-

Table E3: Cross-party redirection decomposed by sector, geographic concentration, and committee-winner ideological distance (pooled across eras).

rdrobust point estimate on `gave_to_winner_t2` toward the focal winner; race-clustered HC1 standard error in parentheses; MSE-optimal bandwidth selected per cell; triangular kernel; local-linear. The within-party estimate is retained as the decomposition reference. Industry is the regulated (utilities/finance/telecom/defense) vs non-regulated regex classification of Appendix Section B.2; geographic concentration is the `pre_state_share` tercile; ideological distance is the committee-winner absolute CFscore distance terciled within the cross-party subsample.

| Cut | Subgroup | Estimate (SE) |
|---|---|---|
| Reference | Cross-party (all) | +0.452 (0.024) |
| | Within-party (all) | +0.276 (0.034) |
| Sector | Regulated (utilities/finance/telecom/defense) | +0.381 (0.040) |
| | utilities | +0.372 (0.041) |
| | finance | +0.343 (0.052) |
| | telecommunications | +0.416 (0.063) |
| | defense | +0.444 (0.058) |
| | Non-regulated | +0.456 (0.022) |
| Geographic concentration | Low tercile | +0.410 (0.035) |
| | Mid tercile | +0.450 (0.029) |
| | High tercile | +0.474 (0.026) |
| Ideological distance | Near (median 0.72) | +0.383 (0.050) |
| | Mid (median 1.13) | +0.344 (0.036) |
| | Far (median 1.76) | +0.337 (0.036) |

ing what body Section 4.5 reports. The motivation is theoretical: in the allocation model (body Section 2) ideological proximity pulls a committee away from a cross-party winner, so a positive cross-party discontinuity isolates the office term $V_O$. The decomposition asks whose office-seeking generates it.

Differences between subgroups, treating the disjoint subsamples as independent: cross minus within +0.176 (standard error 0.041); regulated minus non-regulated $-0.075$ (0.046); geographic high minus low +0.064 (0.043); ideological far minus near $-0.046$ (0.061). Every one of the fourteen subgroup estimates is at least six standard errors from zero; the least precise is the finance sub-industry, +0.343 with a standard error of 0.052.

The cross-party component is both larger than the within-party component (+0.452 versus +0.276) and, by the model mapping, the cleaner office-following signal. Its decomposition by moderator shows the office-following is broad-based. Committees in classically regulated industries redirect at +0.381, statistically indistinguishable from the

+0.456 of the rest of the committee universe (a difference of $-0.075$ with a standard error of 0.046), and the four regulated sub-industries bracket the non-regulated estimate rather than exceeding it; cross-party office-following is a general property of organized giving, not a regulated-sector phenomenon. Geographic concentration adds a gradient in the access-predicted direction (high-concentration +0.474 versus low +0.410), but the gradient is modest (+0.064, standard error 0.043) and, as shown below, does not survive a fully flexible specification, so the paper rests no claim on it. A two-way sector-by-geography cut confirms the gradient is industry-independent: it holds within both regulated and non-regulated committees, and the largest cell is non-regulated committees at high concentration (+0.472), not regulated committees at high concentration (+0.440), so the gradient is not a regulated-sector effect in disguise.

The ideological-distance decomposition ties the pattern to the object the model identifies. Sorting cross-party dyads by the absolute CFscore distance between committee and winner, redirection is +0.383 in the nearest tercile (median distance 0.72), +0.344 in the middle (median 1.13), and +0.337 in the farthest (median 1.76); the far-minus-near difference is $-0.046$ with a standard error of 0.061, not distinguishable from zero. The farthest tercile (committees funding a winner nearly two CFscore units away, across most of the partisan spectrum) redirects at +0.337 (standard error 0.036): office-seeking overcomes a large ideological cost, and on this margin the cost leaves no detectable trace, which is the access-dominant region the model identifies (office weight $V_O$ dominant). The positive proximity weight $V_I$ the model retains is identified by the substitution test of Appendix Section H and the one-party portfolios of Appendix Section P, not by this gradient. The robustly positive redirection among the most opposed dyads also rules out the reading that cross-party redirection is an artifact of committees whose ideological sign is barely defined.

Two methodological points bound the reading. First, the subgroup tercile rdrobust estimates (each in the headline sign convention and directly comparable to the +0.452

cross-party reference) are the primary evidence; the fully-interacted continuous specification `loss × moderator × rv` (kernel-weighted, race-clustered) is used only to check whether a tercile gradient survives a flexible functional form. On that check the sector result is confirmed (the regulated interaction, +0.066, is significant at the five percent level and marginally reduces redirection), the geography gradient does not survive (the loss-by-state-share interaction, +0.032, is not distinguishable from zero), and the distance gradient is absent in both forms (the loss-by-distance interaction is $-0.006$ with a standard error of 0.045); the rv-slope-restricted interaction used elsewhere in the project is avoided here because it manufactures a spurious geography interaction that vanishes once the running-variable slope is allowed to vary by the moderator. Second, the subgroup-difference z-tests treat disjoint subsamples as independent; shared races across terciles make them mildly anti-conservative, so the cross-minus-within contrast, at 4.3 standard errors, holds under that correction, while the geographic high-minus-low and far-minus-near contrasts, at 1.5 and 0.7, do not. The meaning of "cross-party redirection" is inherited from the within-vs-cross decomposition of E.1; this subsection splits that same quantity by moderator and does not re-derive its identification.

# F PAC-side ideology channel

The PAC-side ideological-distance moderator characterizes a layered ideology channel that runs alongside the dominant access redirection. I classify PAC-winner pairs by tercile-tercile match in CFscore space (using `contributor.cfscore` from the DIME committee master file, which is non-missing for every committee there and joins to 98.6 percent of the panel) and by continuous distance. The moderator conditions on a winner trait and is exposed to the Marshall (2024) compensating-differentials concern, so its magnitude is only partially identified; I read this test as a substantively sharper but partially identified companion to the cleanly identified race-side test of Section G. It speaks

to the stronger ideology prediction (whether a PAC's own ideological cost moderates its redirection), while the race-side test cleanly rejects the weaker one.

**Tercile match-distance**

| Match | Estimate | SE | p-value |
|---|---|---|---|
| Same tercile | 0.368 | 0.031 | < 10^-33 |
| Adjacent tercile | 0.330 | 0.023 | < 10^-46 |
| Opposite tercile | 0.211 | 0.041 | 1.2e-6 |

**Continuous interaction within MSE-optimal bandwidth (race-cluster SE)**

`loss × |pac_cfscore − winner_cfscore|` on `gave_to_winner_t2`: unbounded point estimate $-0.226$ (standard error 0.013; 95 percent confidence interval $-0.251$ to $-0.200$). However, this PAC-side interaction conditions on a winner trait and is exposed to the Marshall (2024) compensating-differentials concern; subsection F.1 below reports and applies the partial-identification bounds of Marshall (2024, Section 6) before interpreting the magnitude.

## F.1 Partial-identification bounds following Marshall (2024)

Marshall (2024, Section 6) shows that politician-characteristic regression discontinuity heterogeneity coefficients are biased by compensating differentials at the cutoff: among close races where the winner is ideologically far from the contributing PAC, unobserved race characteristics (candidate quality, district shocks, opponent weakness) systematically differ from those in close races where the winner is ideologically aligned. The estimator therefore conflates the heterogeneity of interest with these compensating differentials. Under the Cauchy-Schwarz-style assumption that the conditional bias satisfies $|\text{Cov}(W, Z \mid rv \approx 0)| \leq K \cdot \text{Var}(Z \mid rv \approx 0)$ for some scalar K characterizing the strength of the unobserved channel, the true coefficient is bounded by

$$[\hat{\beta} - K \cdot \sigma_Z, \;\; \hat{\beta} + K \cdot \sigma_Z]$$

where $\sigma_Z$ is the local standard deviation of the conditioning trait at the cutoff. I compute $\sigma_Z$ with kernel weights at the same MSE-optimal bandwidth used to estimate $\hat{\beta}$.

| K | $\sigma_Z$ | $[\beta_{\text{lo}}, \beta_{\text{hi}}]$ | 95% CI envelope | Excludes 0? |
|---|---|---|---|---|
| 0.5 | 0.9195 | [-0.6857, +0.2338] | [-0.7112, +0.2593] | No |
| 1.0 | 0.9195 | [-1.1454, +0.6935] | [-1.1709, +0.7190] | No |
| 2.0 | 0.9195 | [-2.0649, +1.6130] | [-2.0904, +1.6385] | No |

The bounds bracket zero even at K = 0.5 (the most conservative bias allowance). The point estimate $\beta = -0.226$ is therefore not point-identified under Marshall's partial-identification framework, and I do not interpret its magnitude as a quantified ideology channel. The tercile-match pattern (same / adjacent / opposite) and the 9-cell decomposition above are retained as evidence of a layered ideology channel, real in sign but bounded in magnitude, and the body's language is calibrated accordingly. The PAC-side test speaks to the substantively stronger ideology prediction but only partially identifies it, so the access-dominant reading rests on the cleanly identified race-side test (Section G) and the bounded PAC-side test jointly rather than on either alone.

**Full 9-cell decomposition** (`gave_to_winner_t2`, PAC tercile × winner tercile)

| | WIN_T1 (left) | WIN_T2 (mid) | WIN_T3 (right) |
|---|---|---|---|
| **PAC_T1 (left)** | **+0.414** | +0.501 | +0.209 |
| **PAC_T2 (mid)** | +0.347 | +0.412 | +0.220 |
| **PAC_T3 (right)** | +0.221 | +0.398 | **+0.308** |

Diagonal (same-tercile) cells exceed opposite-tercile cells by 1.4 to 2 times on both extreme-PAC rows: PAC_T1 with WIN_T1 (+0.414) > PAC_T1 with WIN_T3 (+0.209); PAC_T3 with WIN_T3 (+0.308) > PAC_T3 with WIN_T1 (+0.221).

**Why the binary spec disagrees with tercile/continuous.** A binary sign-match parametrization (`compatible` if `sign(pac_cfscore) = sign(winner_cfscore)`) collapses ideology magnitude to direction. Committees cluster heavily at moderate-R (median pac_cfscore $\approx$ +0.35; PAC_T2 holds 65 percent of all panel rows). Binary "compatible" is dominated by weak-R PAC $\times$ weak-R winner pairs, where the roughly +0.30 baseline redirection drives the average. Binary "incompatible" is the small set of D-leaning PACs facing R winners (or vice versa) where selection into the off-sign group correlates with PAC types whose access incentive is larger, producing an aggregate gap (+0.448 vs +0.280) that points opposite the tercile and continuous specifications. Tercile and continuous moderators preserve magnitude information and consistently show distance attenuates redirection.

The tercile and continuous specifications on CFscore are the preferred *PAC-side* parametrizations relative to the binary alternative. The cleanly identified adjudication of access against ideology is the *race-side* test in Section G, which tests the weaker of the two ideology predictions but is least exposed to the Marshall critique; this PAC-side test is its substantively sharper, partially identified companion.

# G Race-level CFscore-gap heterogeneity

Per-quartile rdrobust on `gave_to_winner_t2`:

| Gap quartile | Median gap | Estimate | SE | 95% CI |
|---|---|---|---|---|
| Q1 (low gap) | 1.04 | 0.395 | 0.034 | [0.342, 0.475] |
| Q2 | 1.61 | 0.408 | 0.030 | [0.360, 0.476] |
| Q3 | 1.91 | 0.380 | 0.029 | [0.328, 0.440] |
| Q4 (high gap) | 2.23 | 0.356 | 0.025 | [0.303, 0.400] |

Per-quartile rdrobust on `gave_to_loser_t2`:

| Gap quartile | Estimate | SE | 95% CI |
|---|---|---|---|
| Q1 | -0.269 | 0.036 | [-0.356, -0.216] |
| Q2 | -0.183 | 0.039 | [-0.274, -0.120] |
| Q3 | -0.234 | 0.041 | [-0.331, -0.169] |
| Q4 | -0.137 | 0.032 | [-0.207, -0.083] |

Continuous loss $\times$ CFscore-gap interaction (within the MSE-optimal bandwidth, race-clustered standard errors): `gave_to_winner_t2` $-0.019$ (standard error 0.014); `gave_to_loser_t2` $-0.073$ (0.016); `delta_log_total_giving` $-0.042$ (0.070); `delta_dem_share` (raw) $-0.001$ (0.006); `delta_toward_winner` (sign-aligned, manuscript-primary) +0.012 (0.005). The loser-outcome interaction is four and a half standard errors from zero; the winnerward-composition interaction is small (just over two standard errors), and the other three are within one and a half standard errors of zero.

The loser-loyalty interaction is the only outcome with a meaningful interaction; it is sign-preserving (PACs still drop the loser, just slightly less when the race was ideologically polarized) and does not change the access verdict on the winner side.

# H Ideology-alternative shadow effect

Outcomes are constructed in cycle t+2 as indicators that the PAC made any positive contribution to a non-focal House candidate within $\varepsilon$ CFscore units of the focal loser's CFscore, for four threshold values, plus an intensive count and a geographic complement. `gave_to_ideo_alt_e10_t2` uses $\varepsilon$ = 1.00, a strict subset of the tight threshold; `gave_to_ideo_alt_tight_t2` uses $\varepsilon$ = 1.38, fixed at the 25th percentile of the race-level winner-loser CFscore gap as computed on the 1990–2020 panel (the extended panel's 25th percentile is 1.40) and reported as the manuscript-primary tight threshold; `gave_to_ideo_alt_e15_t2` uses $\varepsilon$ = 1.50, intermediate between tight and

median; `gave_to_ideo_alt_median_t2` uses $\varepsilon = 1.75$, the race-level median CFscore gap; `n_ideo_alt_tight_t2` counts the contributions meeting the tight-threshold criterion (intensive margin); and `gave_to_focal_state_winner_t2` indicates any contribution to a non-focal House winner in the focal state (geographic complement).

| Outcome | rdrobust point | SE | 95% CI | h (pp) | N_h | Implied loss=1 effect |
|---|---|---|---|---|---|---|
| gave_to_ideo_alt_e10_t2 ($\varepsilon$=1.00) | -0.231 | 0.017 | [-0.260, -0.194] | 17.9 | 813,374 | +0.227 |
| gave_to_ideo_alt_tight_t2 ($\varepsilon$=1.38) | -0.161 | 0.017 | [-0.190, -0.124] | 16.1 | 747,821 | +0.158 |
| gave_to_ideo_alt_e15_t2 ($\varepsilon$=1.50) | -0.124 | 0.016 | [-0.151, -0.087] | 15.4 | 723,383 | +0.121 |
| gave_to_ideo_alt_median_t2 ($\varepsilon$=1.75) | -0.063 | 0.016 | [-0.090, -0.029] | 14.9 | 699,005 | +0.062 |
| n_ideo_alt_tight_t2 (count) | -26.54 | 2.54 | [-32.0, -22.06] | 11.6 | 556,320 | +26.7 |
| gave_to_focal_state_winner_t2 (geo complement) | +0.011 | 0.022 | [-0.032, +0.054] | 17.6 | 805,812 | +0.011 (null) |

Sign convention: rdrobust returns the winner-side right limit minus the loser-side left limit, so a negative coefficient means that committees whose pick lost give more on this outcome, and the ideology-replacement channel is positive in loss terms.

Sensitivity to the $\varepsilon$ threshold: the magnitude of the discontinuity shrinks smoothly and monotonically as $\varepsilon$ is loosened (+0.227 at $\varepsilon$ = 1.00, +0.158 at $\varepsilon$ = 1.38, +0.121 at $\varepsilon$ = 1.50, +0.062 at $\varepsilon$ = 1.75), and every threshold preserves the negative sign and statistical significance at conventional levels (each at least three standard errors from zero, the two tightest more than eight). The monotonic shape reflects the mechanical fact that loosening

the similarity threshold incorporates more candidates whose CFscores are further from the focal loser's, attenuating the discontinuity toward the broader-redirection baseline. The manuscript-primary $\varepsilon = 1.38$ (25th-percentile cutoff) lies in the middle of a smooth sensitivity curve rather than at an extremum; the broader $\varepsilon = 1.75$ reading reported in the same body table reinforces the same channel at the median cutoff. The headline ideology-replacement reading is robust to the choice of similarity threshold.

Saturation diagnostic: the binary ideology-alt outcomes saturate near the cutoff (mean rising from about 0.75 at $\varepsilon = 1.00$ to about 0.82 at $\varepsilon = 1.75$), confirming that replacement options are widely available across thresholds. The intensive count margin (`n_ideo_alt_tight_t2`) is the cleanest detection of the ideology-shadow channel at the tight threshold and is reported separately.

# I Geographic-concentration heterogeneity

The moderator `pre_state_share` is the share of a PAC's lifetime House giving through cycle t that went to the focal state. Terciles are cut within each focal cycle; their medians are 0.012 (low), 0.042 (mid), and 0.230 (high), so the distribution is heavily right-skewed.

Per-tercile rdrobust on `gave_to_winner_t2`:

| Tercile | Estimate | SE | 95% CI | h (pp) |
|---|---|---|---|---|
| Low | 0.375 | 0.024 | [0.328, 0.423] | 21.8 |
| Mid | 0.395 | 0.016 | [0.371, 0.431] | 17.9 |
| High | 0.361 | 0.019 | [0.331, 0.406] | 12.4 |

Continuous loss $\times$ pre_state_share OLS interactions (within the MSE-optimal bandwidth, race-clustered standard errors):

| Outcome | Interaction coefficient | SE | p-value | 95% CI |
|---|---|---|---|---|
| gave_to_winner_t2 | +0.137 | 0.009 | 2.0e-52 | [+0.119, +0.154] |
| gave_to_loser_t2 | -0.057 | 0.013 | 4.6e-6 | [-0.082, -0.033] |
| delta_log_total_giving | -0.778 | 0.080 | 2.4e-22 | [-0.935, -0.622] |
| delta_dem_share (raw) | -0.001 | 0.009 | 0.90 | [-0.018, +0.016] |
| delta_toward_winner (sign-aligned, manuscript-primary) | +0.075 | 0.007 | 2.3e-28 | [+0.062, +0.089] |
| gave_to_ideo_alt_tight_t2 | -0.058 | 0.011 | 4.7e-7 | [-0.080, -0.035] |

Because `loss` is the low-`gave_to_winner` side, the discontinuity at moderator value $m$ is $\tau(m) = -(\beta_{loss} + \beta_{loss \times share} \cdot m)$, so a *positive* interaction coefficient shrinks the win-minus-loss gap. Correctly decoded, the three binary-outcome interactions say the same thing: higher geographic concentration attenuates every discontinuity: the winner-redirection gap falls from +0.396 at the low-tercile median to +0.365 at the high-tercile median and +0.329 at pre_state_share = 0.50, the loser-loyalty drop is shallower, and the ideology-alternative substitution weaker. This matches the essentially flat per-tercile estimates above and tracks the narrower giving portfolios of locally concentrated committees (the moderator main effect on winner-giving is $-0.19$), the same breadth mechanic as the single-race attenuation in Table B3 and the donor-size gradient in body Section 4.3, not a district-specific amplification of access. (The `delta_toward_winner` moderated magnitude does rise with concentration, but mechanically: the more concentrated a committee's giving in the focal state, the larger the share of its portfolio the focal race moves.)

Table J1: McCrary density test at the cutoff.

| Test | T-statistic | p-value | density (left) | density (right) |
|---|---|---|---|---|
| McCrary (rddensity local-quadratic) | 9.22 | 2.9e-20 | 0.0158 | 0.0180 |

A fully-interacted verification (`loss*share*rv`, allowing the running-variable slope to differ by moderator level, the specification under which the cross-party geography gradient of Section E.2 vanishes) leaves all three binary-outcome interactions intact: gave_to_winner_t2 +0.142 against +0.137 restricted (more than seven standard errors from zero), gave_to_loser_t2 $-0.089$ and gave_to_ideo_alt_tight_t2 $-0.085$ (each nearly four standard errors from zero). The interaction is therefore statistically robust; what it supports is the uniform-attenuation reading above, and the body accordingly reports the geography moderator as evidence *against* district-specific concentration of the redirection.

# J Validity: McCrary density and bandwidth-stability summary

Table J1 reports the McCrary (2008) density test.

The test rejects continuity, but for a mechanical reason. The running variable is drawn once per race, while the unit of analysis is the committee-race pair: every committee that backed either candidate inherits the same draw, and winners attract more committees than losers, so the pair-level density breaks at the cutoff even if no candidate manipulates any race. McCrary's test targets manipulation by the units that generate the running variable (McCrary, 2008); those units are races, and the race-level density is continuous (subsection J.2). What identification requires is continuity of the conditional expectations of potential outcomes at the cutoff, and the direct falsification check for that is pre-treatment

covariate balance (de la Cuesta and Imai, 2016).

Bandwidth stability for the four co-primary outcomes is reported in the specification curve (Section B). The `gave_to_winner_t2` estimate ranges from 0.34 to 0.39 across bandwidths from 1 to 20 points.

Covariate balance (Table 2 in the body): all six baseline covariates have absolute t-statistics below 2 at the cutoff, five of them below 1; the largest is 1.31, on the number of transactions.

## J.1 Inner-donut covariate balance

This replication-structure defense of the density rejection rests on the fact that the running variable's sign is assigned at the race level and replicated across the committees backing each side, so pair-level density continuity is not implied by the absence of candidate-side sorting; the validity check that bears on identification is covariate balance. To tighten this defense against the methodologist concern that any precise manipulation would concentrate in the narrowest possible window around the cutoff, I re-run the six-covariate balance test on the inner-donut subsample with $|rv| < 0.5$ percentage points. The subsample contains 29,175 PAC-race observations across 10,773 unique PACs, 68 close races, and 16 cycles. Inside the donut the running variable carries negligible information, so the balance test simplifies from the rdrobust local-linear specification of Table 2 in the body to a linear-probability regression of each covariate on the loss indicator with race-clustered HC1 standard errors. The `loss`-coefficient and its t-statistic are the inner-donut balance statistic per covariate.

All six covariates remain balanced at the conventional $|t| < 2$ threshold inside the inner-donut window. The largest absolute t-statistic, 1.41, is on `n_transactions_t`, with `dem_share_t` close behind (1.37); the 68-race subsample is under two percent of the full panel's race count, and both statistics are consistent with sampling variation rather than precise manipulation. Two of the six covariates additionally meet the stricter $|t|$

Table J2: Inner-donut covariate balance ($|rv| < 0.5$ percentage points, 29,175 PAC-race observations across 68 close races).

`diff` is the loss-side minus winner-side covariate mean; SE is race-clustered HC1.

| Covariate | Mean (winner side) | Mean (loser side) | Diff (loser − winner) | SE | $\|t\|$ | p | Balanced ($\|t\| < 2$) |
|---|---|---|---|---|---|---|---|
| log_contribution_to_focal_candidate_t2 | 7.630 | 7.591 | −0.039 | 0.057 | 0.69 | 0.49 | yes |
| n_transactions_t | 2.120 | 1.966 | −0.154 | 0.109 | 1.41 | 0.16 | yes |
| log_total_house_giving_t | 11.333 | 11.155 | −0.178 | 0.169 | 1.05 | 0.30 | yes |
| dem_share_t | 0.434 | 0.550 | +0.116 | 0.085 | 1.37 | 0.18 | yes |
| pac_age_in_dime | 10.712 | 10.163 | −0.549 | 0.700 | 0.78 | 0.44 | yes |
| focal_candidate_cfscore | 0.330 | 0.024 | −0.306 | 0.245 | 1.25 | 0.22 | yes |

< 1 threshold; the four that exceed $|t| = 1$ carry the largest sampling variability in this narrow window, and none approaches conventional statistical significance. The inner-donut balance result reinforces the body's reading that the McCrary density asymmetry is mechanical rather than evidence of precise manipulation at the threshold.

## J.2 Differential PAC attrition at t+2

Three of the four headline outcomes treat dissolved PACs as zero (`gave_to_winner_t2`, `gave_to_loser_t2`, `delta_log_total_giving`); the fourth, `delta_dem_share` / `delta_toward_winner`, is undefined for them (Appendix A.3). If PAC dissolution probability is differential at the cutoff (for example, if PACs whose pick lost are more likely to wind down or merge), all four headline estimates would be biased. I test the concern directly by constructing the binary indicator `active_t2` (equal to 1 if the PAC makes any contribution at cycle t+2 anywhere in DIME, and zero otherwise) and re-running rdrobust with the same triangular kernel, MSE-optimal bandwidth, local-linear polynomial, and race-clustered HC1 standard errors used for the headline outcomes. The discontinuity is +0.0114 (standard error 0.0107; 95 percent confidence interval −0.0096 to +0.0323; MSE-optimal bandwidth 16.28 percentage points; about 751,000 observations within the bandwidth). Committees whose pick narrowly won are about 1.1 percentage points more likely to remain active at t+2, a differential not distinguishable from zero.

Its magnitude stays far below the conventional 0.05 materiality threshold (the confidence interval bounds the differential below 3.2 percentage points, against headline discontinuities of 23 to 38 points), and the raw 5.6-percentage-point loss-side versus winner-side gap in unconditional means (winner-side 0.873, loss-side 0.817) collapses by an order of magnitude at the cutoff. Because the point estimate is positive on the winner side, forced zeros are marginally more common on the loss side; that direction would slightly attenuate the loser-loyalty drop and slightly inflate the winner discontinuity, by at most the 3.2-point bound. The zero-treated-as-zero convention for the four headline outcomes therefore remains immaterial for the paper's conclusions.

## J.3 Pre-trend balance for differenced outcomes

Table 2 in the body reports level-balance tests on six baseline covariates. Two of the four co-primary outcomes are themselves cycle-to-cycle differences (`delta_log_total_giving` and the partisan-composition outcome, reported as `delta_dem_share` in raw form and `delta_toward_winner` in the manuscript-primary sign-aligned form), and for these the appropriate companion test is whether the pre-period version of each delta is balanced at the cutoff: a pre-period differential trend would imply that the post-period estimate mixes a focal-race effect with the continuation of pre-existing trajectories. Constructing the pre-period deltas by panel-internal lookup of each PAC's prior-cycle giving (80.7 percent coverage at t-2; PACs without a focal race at t-2 are excluded from this check) and re-running rdrobust on each:

| Pre-period outcome | Estimate | SE | p | Balanced |
|---|---|---|---|---|
| delta_log_total_giving_pre | +0.023 | 0.038 | 0.33 | yes |
| delta_toward_winner_pre (sign-aligned, manuscript-primary) | +0.027 | 0.007 | $1 \times 10$^-4 | no |
| delta_dem_share_pre (raw) | −0.017 | 0.007 | 0.007 | no |

The total-giving delta is balanced, which strengthens the no-expansion reading: no differential drift in giving levels precedes the cutoff, and none follows it. The winnerward delta is not balanced, and the imbalance is what the design predicts rather than a failure of it. The running variable is the committee's own signed margin, so the two sides of the cutoff are different committees: within a single close race, the winner's backers sit just above zero and the loser's backers just below, leaning toward opposite parties. Close-election assignment randomizes which side won, not which committees backed whom. Two forces then separate the pre-period deltas at the cutoff. First, a share-composition effect: `dem_share` includes the committee's own focal-race contribution, which every committee in the panel made by construction, and orienting the share toward the focal winner's party counts that contribution positively on the winning side and negatively on the losing side. Second, a genuine differential trend: committee giving sorted toward its own party over the sample period, and own-party sorting registers as winnerward drift on the winning side and anti-winnerward drift on the losing side.

To separate the two components I rebuild each committee's House portfolio at cycles t, t-2, and t-4 directly from the contribution records (the reconstruction reproduces the panel's `dem_share_t` exactly), remove the two focal candidates from the share at every endpoint, and re-estimate the pre-period discontinuity; the same construction runs on the individual-donor panel of Appendix Section P:

| Pre-period window | Committees $\beta$ (SE) | Individual donors $\beta$ (SE) |
|---|---|---|
| Focal-inclusive, t-2 to t | +0.0265 (0.0067) | +0.0226 (0.0032) |
| Focal-excluded, t-2 to t | +0.0165 (0.0067) | +0.0083 (0.0016) |
| Focal-excluded, t-4 to t-2 | +0.0233 (0.0066) | +0.0096 (0.0013) |

The share-composition component is 38 percent of the committee wedge and 63 percent of the individual one (individuals concentrate their giving on fewer candidates, so the focal race is a larger slice of the share). The remainder is a differential trend, and the earlier window is the decisive row: it ends two cycles before the focal election and contains no focal-candidate money at either endpoint, yet shows the same widening, so the trend is a property of the two donor groups rather than anything the focal race produced. The trend is not an artifact of the estimator or of the closest races: placebo cutoffs at rv = $\pm$5 and $\pm$10 return nothing (none of the four placebo estimates is more than 1.2 standard errors from zero), donut holes of 0.5 and 1.0 percentage points raise rather than lower it, and fixed bandwidths of 5, 10, and 20 leave it between +0.015 and +0.027. The placebo cutoffs rule out mistaking a smooth gradient in the running variable for a jump; they cannot test the orientation convention itself, which changes sign only at zero by construction.

The estimates the body reports net this trend out. The composite outcome subtracts each committee's pre-period delta from its post-period delta on a matching convention (focal-excluded pre from focal-excluded post, and likewise focal-inclusive), and the discontinuity in the composite is the trend-adjusted shift. The adjustment is identified under the assumption that the differential trend would have continued one window forward at its measured rate; the assumption is not directly testable, but the trend's stability across the two independent pre-period windows (+0.0165 and +0.0233) is the evidence the extrapolation needs. On the focal-excluded convention the trend-adjusted pooled shift is +0.0681 (standard error 0.0100; bandwidth 13.1 percentage points; 76 percent of panel rows enter the composite); on the focal-inclusive convention it is +0.0867 (standard error 0.0107). Because the pre-period drift and the post-period shift have opposite signs on

the winnerward scale, the adjustment enlarges the estimate: the losing side was moving away from the eventual winner's party before the race and toward it after.

### J.4 Within-pool trend-adjusted estimates

The pooled shift could in principle be one pool's movement. Splitting by the winner's party, with every estimate on the focal-excluded convention:

| Pool | Post (SE) | Pre-trend (SE) | Trend-adjusted (SE) | 95-percent CI |
|---|---|---|---|---|
| Pooled | +0.0497 (0.0074) | +0.0165 (0.0067) | **+0.0681 (0.0100)** | [+0.047, +0.086] |
| D-winner | +0.0430 (0.0106) | +0.0035 (0.0102) | **+0.0427 (0.0152)** | [+0.007, +0.067] |
| R-winner | +0.0538 (0.0100) | +0.0236 (0.0083) | **+0.0820 (0.0128)** | [+0.060, +0.110] |

Both pools move toward the winner: PACs shift toward the winning Democrat when a Democrat wins and toward the winning Republican when a Republican wins, so co-partisan migration cannot generate the pooled estimate. The differential trend is not symmetric across pools. Most of it comes from the R-winner pool (+0.0236 in the adjacent window, +0.0250 in the earlier one), so the R-winner adjustment is large; the D-winner trend is near zero in the adjacent window and +0.0205 in the earlier one, so its trend-adjusted estimate is the less settled of the two, and its raw and adjusted values coincide (+0.043). Under either reading the direction is the same in both pools, and the R-winner estimate, where the trend is measured most stably, adjusts upward rather than downward: the raw within-pool magnitudes are conservative.

## K Multiple-testing correction

Bonferroni and Benjamini-Hochberg adjustments for the four-outcome family applied to the four pre-specified outcomes:

| Outcome | p (raw) | p (Bonferroni) | p (BH) | Passes Bonferroni | Passes BH |
|---|---|---|---|---|---|
| gave_to_winner_t | 2.2e-128 | 8.9e-128 | 8.9e-128 | yes | yes |
| gave_to_loser_t | 2.3e-24 | 9.0e-24 | 4.5e-24 | yes | yes |
| delta_log_total_giving | 0.90 | 1.00 | 0.90 | no | no |
| delta_dem_share | 0.14 | 0.55 | 0.18 | no | no |

Two of the four pre-specified co-primary outcomes remain significant after correction by many orders of magnitude. The total-giving outcome is a genuine informative null (the exit-robust margins bound any budget expansion below five to seven percent, falsifying the budget-expansion hybrid; Appendix Section Q). The raw partisan-composition outcome (`delta_dem_share`) is non-significant under family-wise correction because the pooled estimate is a cancellation artifact of two opposite-sign within-pool movements rather than a substantive null. Reported on the manuscript-primary winner-ward parametrization (`delta_toward_winner`, the deterministic sign-aligned derivative documented in Section A.3), the partisan-composition outcome is itself a large positive discontinuity (0.0623 in absolute value on the rdrobust convention, standard error 0.0080, nearly eight standard errors from zero and well beyond any family-wise threshold). The corrected pattern reads as three large effects (winner up, loser down, partisan composition winnerward) plus one informative null (total giving unchanged), which together falsify the hybrid and partisan-loyalty alternatives. The winnerward parametrization is reported as a deterministic transform of the pre-specified `delta_dem_share`, not as an additional test, so the Bonferroni and BH thresholds are unchanged. The four outcomes are of course not independent tests (a committee that redirects toward the winner mechanically moves its portfolio composition winnerward), which is why the body reads them as one joint signature rather than four separate confirmations; the decomposition in Appendix Section O quantifies the mechanical component of the composition outcome directly (20.2 percent) and further separates the remainder into an other-contested-race

projection component and a small genuinely portfolio-wide part concentrated in incumbent defeats.

# L Loser-recontest sensitivity (active disengagement vs mechanical absence)

The body's loser-loyalty drop (Section 4.1) potentially conflates two channels: (a) **active disengagement**, where committees stop supporting a defeated candidate who is again on the ballot at t+2, and (b) **mechanical absence**, where the focal loser does not run again and `gave_to_loser_t2 = 0` is forced by construction. Whether a defeated candidate "runs again" is itself measured at three tiers that differ by a factor of two, so this section labels every rate it reports by tier. (i) Filing: the loser has a candidate committee on file for cycle t+2 in DIME's candidacy roster. (ii) Primary: the loser appears in a t+2 primary. (iii) General ballot: the loser is on the t+2 general-election ballot for the same seat, verified against the MIT Election Lab candidate file by last name within state and district. Rates are weighted as the estimator weights them, over committee-race rows on the losing side of the cutoff with the triangular kernel at the loser outcome's bandwidth. At the cutoff, 88 percent of focal losers file for some office and 82 percent file for the same House seat; 52 percent appear in a primary for some office and 46 percent for the same seat; but only 39 percent are on the same seat's general ballot by DIME's own election fields and 32 percent by the MIT-verified roster, and 21 percent face the same opponent again. The filing tier counts candidacies that never reach a ballot: a committee that raises money for a primary the candidate loses or abandons, or that only retires debt. Not filing at all is rarer at the cutoff (13 percent of loss-side rows within seven points) than in lopsided races (46 percent), so the observations that carry the local-linear fit are those in which the loser at least attempted a return.

*Scope.* The recontest analysis is estimated on focal cycles through 2020, where the t+2

candidacy and ballot records are complete (DIME's 2024 candidacy records cover federal offices only); the loser-loyalty drop on this window is $-0.224$ (the full-panel headline is $-0.227$).

## L.1 Four estimates

| Specification | $\beta$ | SE | 95% CI | h (pp) | N (PAC × race) |
|---|---|---|---|---|---|
| (1) Unrestricted (full panel) | $-$**0.2244** | 0.0244 | [$-0.284$, $-0.188$] | 6.79 | 1,674,412 |
| (2) Filed for any office at t+2 | $-0.2425$ | 0.0259 | [$-0.305$, $-0.203$] | 7.00 | 1,038,200 |
| (3) Filed for the same House seat at t+2 | $-0.2304$ | 0.0264 | [$-0.294$, $-0.190$] | 7.26 | 918,155 |
| (4) On the same seat's general ballot at t+2 (MIT-verified) | $-$**0.3956** | 0.0238 | [$-0.455$, $-0.362$] | 8.93 | 222,622 |

Specifications (2) and (3) restrict to committee-race observations whose focal loser appears in DIME's candidacy roster at t+2 (the filing tier), spec (3) for the same state and district; spec (4) restricts to losers verifiably on the same seat's general ballot (762 races). The three filing-tier estimates lie within two percentage points of one another because the filing roster is dominated by candidacies that never reach a ballot, where `gave_to_loser_t2` is near zero on both sides of the cutoff and the rows contribute little to the discontinuity. Removing those rows, spec (4) makes the discontinuity larger in magnitude, $-0.396$ against $-0.224$: on the losing side, 44 percent of the committees that backed a rerunning loser fund that candidate again, against 4 percent of winner-side

committees funding the losing opponent (side limits at the spec-4 bandwidth).

## L.2 Candidate-level premiums by rematch status

The office-level outcome cannot say whether that 44 percent is loyalty sustained or loyalty withdrawn, because it compares giving to two different candidates. Table L2 re-estimates the candidate-level premiums of Section 4.1 (the discontinuities in `gave_to_pick_t2` and `gave_to_opp_t2`, Appendix Section B.4) within the ballot-verified rematch subsample and within its complement, the races whose loser is not on the same seat's general ballot, with the side limits at each cell's MSE-optimal bandwidth.

Table L2: Candidate-level premiums by whether the defeated candidate is on the same seat's t+2 general ballot (focal cycles through 2020).

rdrobust conventional local-linear estimates with robust bias-corrected standard errors, triangular kernel, MSE-optimal bandwidth per cell, race-clustered SEs; the crossover premium is reported loss-side over win-side. Side limits are the kernel-weighted local-linear intercepts on each side at the cell's bandwidth.

| Subsample | Own-pick premium (SE) | Own pick funded: win side / loss side | Crossover premium (SE) | Opponent funded: win side / loss side | Races |
|---|---|---|---|---|---|
| Loser on the same seat's ballot | +0.008 (0.044) | 0.445 / 0.438 | +0.006 (0.013) | 0.045 / 0.051 | 762 |
| Loser not on that ballot | +0.402 (0.022) | 0.478 / 0.076 | +0.117 (0.011) | 0.010 / 0.127 | 4,849 |
| All races through 2020 | +0.203 (0.032) | 0.454 / 0.251 | +0.069 (0.011) | 0.026 / 0.095 | 5,611 |

Where the defeated candidate is on the same ballot again, both premiums vanish: backers of the narrow loser fund that candidate at the same rate (44 percent) as backers of the narrow winner fund the winner, and neither side's backers cross to the other candidate (5 percent on each side). Where the defeated candidate has left the ballot, the premiums are +40 and +12 points: 8 percent of the loser's backers still give to the candidate, 13 percent give to the winner they had opposed, and 1 percent of the winner's backers give to the departed loser. The Table 3 premiums are the weighted combination of the two, and they are premiums of staying in contention for the seat: office keeps the committee's relationship alive by keeping the candidate in the field, and a committee whose candidate

has left the field moves to the officeholder rather than to nobody. A committee whose candidate is contesting the seat again neither abandons that candidate nor hedges by funding the incumbent.

### L.3 Implication and caveat

The population-level loyalty drop is therefore mostly the end of candidacies rather than the abandonment of active ones: about two-thirds of narrow losers are not on the same seat's general ballot at t+2, and it is in those races that the loyalty drop and the crossover premium are paid. This does not weaken the access reading of Section 4.1, which turns on where the money goes once the relationship has no candidate, and it sharpens it: the money goes to the officeholder. One caveat attaches. The split conditions on the loser's decision to run again, which is made after the outcome and may itself respond to backers' loyalty (a candidate whose committees stay is likelier to run), so the rematch subsample is selected on the loss side toward loyal backers and the 44 percent is an upper bound on conditional loyalty; the complement is selected the opposite way. The unconditional premiums of Table 3 are the causal objects; the split describes the channel through which they arise.

Scripts: `code/16b_loser_recontest_sameseat.R` (specifications 1–4), `code/69_recontest_` (tiers), `code/45_recontest_window_audit.R` (filing rates by window), `code/75_rematch_tier` (Table L2).

## M Within-PAC fixed-effects sensitivity

The headline estimates in Section 4 use a regression-discontinuity design that identifies a between-PAC contrast at the cutoff: the difference in average behavior between PACs whose pick barely won and PACs whose pick barely lost. The substantive language of "redirection" frames this contrast as a within-PAC counterfactual, in which a single PAC

Table M1: Within-PAC fixed-effects sensitivity at MSE-optimal bandwidth.

$\beta_{\text{loss}}$ is the parametric LPM coefficient on the loss indicator with PAC and cycle fixed effects absorbed; the rdrobust loss-effect column reports the corresponding rdrobust estimate at the same fixed bandwidth, re-signed as the equivalent loss=1 effect; distance is reported in units of the rdrobust standard error.

| Outcome | h (pp) | N obs | PACs with both loss states | $\beta_{\text{loss}}$ (within-PAC FE) | SE($\beta_{\text{loss}}$) | rdrobust loss-effect | Distance (SE units) |
|---|---|---|---|---|---|---|---|
| `gave_to_winner_t2` | 15.09 | 710,846 | 19,084 (29.0%) | -0.392 | 0.011 | -0.375 | 1.06 |
| `delta_toward_winner` | 12.35 | 501,092 | 13,282 (54.2%) | +0.061 | 0.006 | +0.062 | 0.11 |

behaves one way when its pick wins and another way when its pick loses. Section 3 of the body acknowledges this distinction explicitly. This appendix supplies the within-PAC fixed-effects sensitivity that validates the prose framing empirically.

I restrict the sample to observations within the extended panel's MSE-optimal bandwidth for each outcome, read directly from the primary specification (h = 15.09 percentage points for `gave_to_winner_t2`, h = 12.35 percentage points for the sign-aligned `delta_toward_winner`), and estimate a parametric local-linear LPM:

$$y_{i,t} = \alpha_i + \gamma_t + \beta \cdot \text{loss}_{i,t} + \delta \cdot \text{rv}_{i,t} + \theta \cdot (\text{loss}_{i,t} \times \text{rv}_{i,t}) + \varepsilon_{i,t},$$

where $\alpha_i$ is a PAC fixed effect, $\gamma_t$ is a cycle fixed effect, $\text{loss}_{i,t}$ is the focal-race-outcome indicator (1 if the PAC's pick lost), and $\text{rv}_{i,t}$ is the running variable (focal-race vote-share margin, signed positive when the PAC's pick is the eventual winner). Standard errors are clustered at the focal-race level. Estimation uses `fixest::feols` to accommodate the high-dimensional PAC fixed effect (approximately 16,400 to 26,300 absorbed levels per outcome after singleton removal).

For both outcomes the within-PAC coefficient is sign-consistent with the rdrobust headline and lies close to it: 1.06 standard errors for the binary headline outcome, 0.11 standard errors for the sign-aligned partisan-composition derivative. Identifying within-PAC variation requires the same PAC to be observed under both loss = 0 and loss = 1

within the bandwidth; this holds for 29.0 percent of PACs in the binary-headline sample and 54.2 percent in the partisan-composition sample, providing a sufficient base for precise estimation. The within-PAC estimate on the binary headline outcome is 35.4 standard errors from zero.

Two notes on interpretation. First, the within-PAC coefficient on the binary headline is slightly larger in absolute value than the between-PAC rdrobust estimate (0.392 versus 0.375). Netting out PAC-specific baselines does not attenuate the discontinuity; the within-PAC contrast at the cutoff is at least as sharp as the between-PAC contrast. Second, the parametric form above is a sensitivity check rather than a primary specification: it does not adopt rdrobust's bias correction and is reported only at the MSE-optimal bandwidth selected for the headline outcome. The close match between the between-PAC and within-PAC estimates indicates that the strategic-redirection prose interpretation is consistent with the within-donor counterfactual the language implies.

## M.1 Within-PAC event-study trajectory

The fixed-effects sensitivity above nets out time-invariant committee heterogeneity but is still a single-period between-outcome contrast within the bandwidth. To recover the within-donor *dynamics* (the evidence body Section 4.6 reports), I track each committee-by-focal-race unit's giving to the two fixed focal candidates, the eventual winner and the eventual loser, across relative election cycles $\tau$ = -2, -1, 0, +1, +2, where $\tau$ = 0 is the focal cycle and $\tau$ = +1 coincides with the paper's t+2 outcome cycle. The event study is estimated on focal cycles through 2020: $\tau$ = +2 for the 2022 cohort is the 2026 cycle, beyond the data, so the estimates below are unchanged by the extension of the panel to 2022. Outcomes are re-streamed from DIME for calendar cycles 1986–2024 and cached; as a validation, giving at $\tau$ = +1 reproduces the canonical `gave_to_winner_t2` and `gave_to_loser_t2` outcomes on this window exactly (agreement 1.0000). The sample is cutoff-local at the pre-extension headline bandwidth (h = 14.29 pp): 228,210 loss-side

Table M2: Within-PAC giving trajectories to the fixed focal candidates (pre-extension headline bandwidth h = 14.29 pp).

Cell entries are the share of committee-by-race units contributing to the named target in the given relative cycle; $\tau$ = +1 is the paper's t+2 outcome cycle.

| Side and target | $\tau$ = -2 | $\tau$ = -1 | $\tau$ = 0 | $\tau$ = +1 (t+2) | $\tau$ = +2 |
|---|---|---|---|---|---|
| Loss side: P(give winner) | 0.015 | 0.018 | 0.028 | 0.103 | 0.070 |
| Loss side: P(give loser) | 0.106 | 0.219 | 1.000 | 0.096 | 0.051 |
| Win side: P(give winner) | 0.179 | 0.323 | 1.000 | 0.478 | 0.253 |
| Win side: P(give loser) | 0.006 | 0.010 | 0.009 | 0.008 | 0.010 |

*Note:* Restricting to units within five percentage points of the cutoff leaves the loss-side winner-following rise essentially unchanged (0.021, 0.024, 0.028, 0.106, 0.067 across the five cycles). The $\tau$ = 0 loss-side give-to-loser value of 1.000 reflects that the focal loser is, by construction, the candidate the committee supported in the focal cycle.

and 391,719 win-side committee-by-race units, of which 17,435 committees (30.0 percent) contribute both a close win and a close loss. Estimation uses `fixest::feols` with unit fixed effects and the pre-campaign cycle $\tau$ = -2 as the reference, since giving to a candidate peaks mechanically during that candidate's own campaign ($\tau$ = -1), which is not a valid pre-treatment baseline.

# N Push-versus-pull adjudication via winner committee chair status

The push-versus-pull distinction within the access view (Section 2 of the body) asks whether the post-focal redirection is initiated by the PAC (push) or by the winner's office staff and party-leadership solicitation infrastructure (pull). Direct adjudication would require contact-level records on which side opened the post-focal relationship, which DIME does not contain. This appendix reports two indirect tests that converge against a pull reading: a decomposition by winner incumbency type (already discussed in Appendix Section B.2 and summarized below for cross-reference) and a direct heterogeneity test using external congressional committee assignment data.

Table M3: Within-donor event-study estimates, loss side (unit fixed effects, reference $\tau$ = -2).

The first two rows are within-unit time profiles for a single target; the third is the winner-versus-loser within-unit difference-in-differences. Coefficients are changes relative to the $\tau$ = -2 baseline.

| Coefficient (vs. $\tau$ = -2) | $\tau$ = -1 | $\tau$ = 0 | $\tau$ = +1 | $\tau$ = +2 |
|---|---|---|---|---|
| P(give winner), within-unit | +0.003 | +0.013 | +0.088 | +0.055 |
| P(give loser), within-unit | +0.113 | +0.894 | -0.010 | -0.055 |
| Winner-vs-loser gap (within-unit DiD) | -0.110 | -0.881 | +0.097 | +0.110 |

*Note:* The $\tau$ = +1 within-unit rise in loss-side giving to the winner (+0.088) has p = 1 $\times$ 10^-89. The winner-minus-loser difference-in-differences (the change in the winner-versus-loser gap relative to $\tau$ = -2) is -0.110 at $\tau$ = -1, +0.097 at $\tau$ = +1 (p = 6 $\times$ 10^-16), and +0.110 at $\tau$ = +2 (p = 4 $\times$ 10^-25); combined with the -0.091 raw gap at the $\tau$ = -2 baseline, the within-committee preference flips from favoring the former pick before the race to favoring the winner after it. On the win side the same difference-in-differences is positive throughout (baseline gap +0.173, rising to about +0.47 at $\tau$ = +1), reflecting incumbency continuation. A naive loss-versus-win difference-in-differences on giving to the winner is confounded, because that outcome is incumbency continuation on the win side but opponent-following on the loss side: it shows a spurious +0.141 win-side pre-trend at $\tau$ = -2 and a $\tau$ = +1 estimate of -0.070, far below the cross-sectional level gap. The clean within-donor object is therefore the within-side winner-versus-loser comparison reported here.

## N.1 Incumbency-type decomposition (summary)

The decomposition by the incumbency type of the focal-race winner, restated from Appendix Section B.2 for cross-reference:

| Winner incumbency type | $\hat{\beta}$ on `gave_to_winner_t2` | Pull prediction |
|---|---|---|
| Open seat (no prior staff or donor pipeline at t) | +0.43 | LOWEST (infrastructure absent) |
| Incumbent reelected (established staff, donor pipeline) | +0.44 | HIGHEST (most-developed pull capacity) |
| Challenger won (new staff at t+1, displaces incumbent) | +0.37 | LOW (no infrastructure inherited) |

The observed estimates (open seat +0.43, incumbent reelected +0.44, challenger won +0.37) do not show the incumbent-reelected advantage the pull mechanism predicts: open-seat winners arrive at t+1 with the least developed donor-relations apparatus, yet they generate redirection indistinguishable from re-elected incumbents' (a difference of 0.007 with a standard error of 0.028; era-adjusted +0.43 against +0.44; Appendix Section B.2 reports the external incumbency classification and the era composition of the subsamples). The pattern is more consistent with PAC-side strategic recalibration than with winner-side solicitation effort, but the open-versus-incumbent contrast alone cannot decide between push and pull; the chair test below is the sharper test. Full discussion is in Appendix Section B.2.

## N.2 Committee chair heterogeneity (Stewart-Woon House assignments)

The incumbency-type decomposition uses an aggregate proxy for winner-side solicitation capacity. A sharper test uses external data on each winner's committee position at the relevant Congress. Standing-committee chairs operate the most developed donor-relations pipelines in Congress: chairs control hearing schedules and bill markups, com-

Table N1: Subgroup rdrobust on `gave_to_winner_t2`, full sample vs by winner chair status (in-coverage subsample, Congresses 103–115).

rdrobust with triangular kernel, MSE-optimal bandwidth, race-clustered HC1 standard errors. Standard errors in parentheses.

| | Winner not chair | Winner is chair |
|---|---|---|
| **Full-committee chair** | | |
| $\hat{\beta}$ (SE) | +0.362*** (0.016) | +0.327*** (0.087) |
| Bandwidth h (pp) | 14.4 | 14.3 |
| N races (within h) | 4,288 | 231 |
| N obs (within h) | 438,692 | 16,269 |
| **Any chair (full + subcmte)** | | |
| $\hat{\beta}$ (SE) | +0.359*** (0.016) | +0.393*** (0.067) |
| Bandwidth h (pp) | 14.7 | 20.0 |
| N races (within h) | 4,237 | 282 |
| N obs (within h) | 445,578 | 29,532 |

mand the largest committee staff, and maintain the most active fundraising operations on behalf of the majority. If the post-focal redirection were materially driven by winner-side solicitation, the chair-led winners would show systematically stronger redirection than rank-and-file members.

**Data source.** Stewart and Woon's House Committee Assignment Data (Congresses 103–115; updated 17 November 2017; http://web.mit.edu/cstewart/www/data/house_assignments_1 115-3.xls) records every member-committee assignment with party-status code (1 = majority original, 2 = minority original) and rank within party status. Chairs are identified as members with party-status = 1 and rank = 1 within a standing committee in a given Congress. I construct two indicators: `is_full_cmte_chair` (chair of any standing committee, codes < 200 in Stewart-Woon's coding) and `is_any_chair` (full or subcommittee chair). The crosswalk from bonica.rid to the real ICPSR number is constructed from the DIME recipients file (DIME's `ICPSR` field is Bonica-encoded as ICPSR $\times$ 10000 + cycle; `ICPSR2` is the real four-to-five-digit ICPSR used by Stewart-Woon and Voteview). Coverage spans Congresses 103–115 (focal cycles 1992–2016, t+2 outcomes 1994–2018), corresponding to 70.9 percent of the full panel; focal cycles 1990 and 2018 through 2022 are missing chair coverage.

Table N2: Continuous loss × chair interaction (parametric LPM at MSE-optimal bandwidth, in-coverage subsample).

`gave_to_winner_t2` regressed on `loss * rv + loss * chair + rv * chair` within $|rv| \leq 14.49$ percentage points (effective N = 459,322; 1,148 races); race-clustered standard errors. The interaction coefficient `loss × chair` is the chair-led increment to the loss=1 effect on the binary headline.

| Interaction | Coefficient | SE | 95-percent CI | p |
|---|---|---|---|---|
| `loss × is_full_cmte_chair` | -0.015 | 0.080 | [-0.172, +0.142] | 0.85 |
| `loss × is_any_chair` | -0.089 | 0.072 | [-0.230, +0.052] | 0.22 |

Neither subgroup contrast is distinguishable from zero: full-committee chairs receive slightly less redirection than non-chairs ($-0.035$, against a chair-subgroup standard error of 0.087), any-chair winners (including subcommittee chairs) slightly more (+0.034, standard error 0.067). The chair-led subgroups are imprecise because the close-election sample contains few chair-led races (231 and 282 within bandwidth), and the subgroup point estimates also absorb the base-rate gap on the winner-side limit (chair-led winners attract very high baseline PAC giving) rather than identifying redirection magnitude per se.

Neither interaction coefficient is distinguishable from zero, and both 95-percent confidence intervals exclude amplification of redirection by winner-side power larger than about half the headline effect (the upper bounds are +0.14 and +0.05 against the +0.375 headline). The full-committee chair interaction is $-0.015$ (the same loss effect for chair-led and non-chair-led winners); the any-chair interaction is $-0.089$ (chair-led winners show a marginally stronger drop, directionally opposite to pull amplification). The continuous LPM is preferred over the subgroup rdrobust for adjudication because the two subgroups are estimated at the same bandwidth and on the same baseline, removing the ceiling-saturation artifact that distorts the subgroup rdrobust on any-chair winners.

The chair test is statistically conservative on the precision side: chair-led winners are a small slice of the close-election sample (committee chairs typically hold safe seats), and the interaction confidence intervals span roughly $\pm 0.15$. A pull mechanism that amplified redirection by 10 percentage points among chair-led winners would lie within the full-

committee interval, and just outside the any-chair one, so it is not reliably detectable here. What the chair test cleanly excludes is a pull mechanism of the magnitude that "winner staff solicitation drives most of the redirection" would imply: such a mechanism would generate interaction coefficients of +0.20 or larger, well outside the estimated confidence intervals.

**Combined verdict.** Both the incumbency-type decomposition (N.1) and the committee chair interaction (N.2) point in the same direction: redirection magnitude is insensitive to plausible variation in winner-side solicitation capacity. The chair coverage limitation (Congresses 103–115; focal cycles 1992–2016) and the sample-size constraint inside the close-election bandwidth leave room for a small pull component the design cannot detect, but the two tests jointly rule out the pull mechanism as the dominant driver of the +0.375 winner-redirection. The framework in Section 2 identifies the push reading as the operative description of the access channel; the empirical evidence reported here supports that reading.

# O Anatomy of the winnerward partisan-composition outcome: mechanical, projected, and portfolio-wide components

The winnerward partisan-composition outcome `delta_toward_winner` is computed on the partisan composition of the PAC's **full House contribution portfolio** in cycle t and cycle t+2. The two focal-race candidates are themselves part of that portfolio, and for the 93 percent of panel rows belonging to committees with picks in more than one focal race that cycle, so are the candidates of the committee's *other* focal races. The pooled discontinuity (+6.23 percentage points at the body Table 3 bandwidth, which the re-selected bandwidth of this decomposition reproduces in Table O1, and the figure the steps below

decompose) therefore has three analytically distinct components: (A) a mechanical component from the PAC's reallocation within the focal race itself; (B) a projection component from the PAC's redirections in its other contested races, mapped onto the focal winner's party axis; and (C) a genuinely portfolio-wide component in races where the committee holds no pick. This appendix separates the three, shows that (B), not portfolio-wide reorientation, accounts for most of the focal-excluded remainder, and locates the genuine component (C) in one configuration: incumbent defeats.

**Construction.** I re-extract the DIME contribution database to recover the exact amount each PAC gave to the focal-race winner (`amt_winner_t`) and focal-race loser (`amt_loser_t`) in cycle t. The cycle-t+2 amounts (`amt_winner_t2`, `amt_loser_t2`) are already in the analytic panel. For each PAC-by-cycle observation, I subtract the focal-race contributions from the numerator and denominator of the PAC's Democratic-share computation, allocating the contribution to the D side or R side according to the focal candidate's party:

$$\text{dem\_share\_excl}_{i,t} = \frac{\text{dem\_giving}_{i,t} - \text{focal\_dem}_{i,t}}{\left(\text{dem\_giving}_{i,t} - \text{focal\_dem}_{i,t}\right) + \left(\text{rep\_giving}_{i,t} - \text{focal\_rep}_{i,t}\right)},$$

where $\text{focal\_dem}_{i,t}$ is the PAC's contribution to whichever of the two focal candidates is a Democrat (zero if both are Republican), and analogously for $\text{focal\_rep}_{i,t}$. The focal-excluded delta is the difference of focal-excluded shares across cycles, and the focal-excluded winnerward outcome `delta_toward_winner_excl` applies the same sign normalization as the primary outcome (positive in D-winning races, negative in R-winning races on the raw delta; sign-aligned to mean "movement toward the focal-race winner's party in the PAC's non-focal House portfolio"). The re-extract recovers exact contributions to both focal candidates in cycle t, including the non-dominant amounts of dual-pick PACs that the primary build's dominant-amount rule discards (19,413

Table O1: Decomposition of the winnerward partisan-composition discontinuity into mechanical and spillover components.

Sign-normalized winnerward outcome `delta_toward_winner` and its focal-excluded counterpart `delta_toward_winner_excl`, rdrobust with triangular kernel and race-clustered standard errors; MSE-optimal bandwidth re-selected for each outcome. Magnitudes reported as $|\hat{\beta}|$ following the sign convention documented for the primary outcome in body Section 3.

| Outcome | $\lvert\hat{\beta}\rvert$ | SE | p (robust) | 95-percent CI | h (pp) | N effective |
|---|---|---|---|---|---|---|
| `delta_toward_winner` (full portfolio) | 0.0623 | 0.0080 | $< 10^{-13}$ | [0.0448, 0.0762] | 12.3 | 501,092 |
| `delta_toward_winner_excl` (focal-race removed) | 0.0497 | 0.0074 | $< 10^{-10}$ | [0.0336, 0.0628] | 13.1 | 527,180 |
| **Mechanical (A) = full − excl** | **0.0126** | — | — | — | — | — |
| **Spillover (B) = excl** | **0.0497** | 0.0074 | $< 10^{-10}$ | [0.0336, 0.0628] | 13.1 | 527,180 |

dual-pick committee-race pairs panel-wide, both amounts recovered exactly).

**Focal-race share of the PAC's portfolio.** The mechanical component is bounded by the share of the PAC's total D and R giving that flowed to the focal-race candidates. Within the MSE-optimal bandwidth window ( | rv | < 13 pp; N = 612,527), the mean focal share is 11.3 percent at cycle t and 2.8 percent at cycle t+2. The asymmetry reflects portfolio diffusion: the focal-cycle contribution is concentrated on the focal candidate by construction, while two cycles later the PAC's giving spreads across many other House races. Even at the higher cycle-t share, focal contributions are a minority of the PAC's House portfolio: the partisan-composition outcome is fundamentally a portfolio measure, not a focal-race measure, regardless of how the decomposition lands.

**Step 1: focal-race exclusion (Table O1).** The mechanical component (A) accounts for 20.2 percent of the full pooled discontinuity (1.26 of 6.23 percentage points); the focal-excluded remainder is +4.97 percentage points (standard error 0.74). On its own, Table O1 would suggest a large portfolio-wide reorientation. The two further steps below show that reading is wrong for multi-race committees.

**Step 2: union exclusion (Table O2).** Excluding the candidates of *all* of the committee's own focal races that cycle (not just the current row's two) collapses the pooled remainder

Table O2: Focal-excluded versus union-excluded winnerward composition discontinuities.

Magnitudes are absolute values of the rdrobust discontinuity, with direction labels; triangular kernel, MSE-optimal bandwidth per cell, race-clustered standard errors.

| Subsample | Focal-excluded | Union-excluded (all own races) | Direction (union) |
|---|---|---|---|
| Pooled | 0.0497 (p < 10^-10) | 0.0088 (p = 0.13) | winnerward |
| Pick was sitting incumbent | 0.0815 (p < 10^-12) | 0.0302 (p < 0.001) | winnerward |
| Pick opposed a sitting incumbent | 0.0163 (p = 0.003) | 0.0297 (p < 10^-5) | own-party |
| No incumbent among top two | 0.0072 (p = 0.42) | 0.0335 (p < 0.001) | own-party |

from +4.97 to +0.88 percentage points (standard error 0.54), not distinguishable from zero. Split by the incumbency configuration of the pick's contest (classification as in the pick-incumbency rows of Section B.3), the genuinely non-focal component survives in exactly one place: committees whose *sitting-incumbent* pick was ousted reorient their no-stake portfolio toward the winning party by +3.02 points (more than three standard errors from zero); after failed challenges and open-seat losses the no-stake portfolio moves *toward the committee's own party* (+3.0 and +3.4 points).

**Step 3: why the gap is projection, not behavior (Table O3).** Two leave-one-out discriminators on multi-race committees separate the candidate explanations. First, the committee's *other* picks' win rate is smooth through the focal cutoff (pooled discontinuity $-0.0020$, less than one standard error from zero, and no configuration beyond one and a half standard errors); the collapse in Step 2 is not driven by correlated outcomes across the committee's races. Second, the committee's winner-following *in its other focal races* is also unchanged at the cutoff (pooled $-0.0031$, a third of a standard error from zero, and no configuration beyond one standard error); losing one race does not intensify redirection elsewhere. With both channels flat, the focal-excluded remainder that Step 2 removes is a **sign-projection artifact**: committees run same-party-tilted slates, so their (unchanged, balanced) redirections in other contested races score as movement *toward* the focal winner's party on the losing side of the cutoff and *away* from it on the winning side, oppo-

Table O3: Leave-one-out discriminators at the focal cutoff (multi-race committees).

| Outcome | Pooled point | p | Reading |
|---|---|---|---|
| Other picks' win rate | -0.0020 | 0.53 | other races' outcomes balanced |
| Winner-following in other focal races | -0.0031 | 0.60 | no cross-race behavioral response |
| Loser-support in other focal races | -0.0015 | 0.14 | no cross-race behavioral response |

site projections of identical behavior. Consistent with the required leverage, giving to other focal candidates amounts to a median 84 percent of the union-excluded denominator among loss-side multi-race committees (computed on the pre-extension panel; the union-exclusion estimates above are re-estimated on the extended panel).

**Substantive implication.** The anatomy recalibrates what the composition outcome establishes. Against partisan loyalty, the facts that matter are the outcome's *sign* (the pooled portfolio moves toward whichever party wins, where a loyal committee whose pick loses should tilt toward its own party) and the cross-party redirection that generates it (Section E.2: the farthest-ideology tercile still redirects at +0.337). What the outcome does *not* establish is a general portfolio-wide party reorientation: outside the committee's own contested races, rebalancing is small in pool and concentrated where the focal outcome genuinely carries news about party strength, the defeat of a sitting incumbent (+3.02 points, more than three standard errors from zero). That pattern is itself informative: an incumbent ousted is a rare, informative event about a party's standing, and committees respond to it portfolio-wide, while an incumbent's routine survival prompts no such updating. The small own-party tilts after failed challenges and open-seat losses are an order of magnitude below the headline discontinuities, and no claim in the body depends on them.

# P Donor-type heterogeneity and the individual-donor comparison

This section supports body Section 4.3. It documents (i) the committee-master crosswalk used to split the committee sample by Federal Election Commission committee type and its coverage, (ii) the redirection by committee type, (iii) the external purpose classification from the Center for Responsive Politics and the anatomy of cross-type conversion differences, and (iv) the parallel individual-donor design, the within-individual donation-size gradient, and the breadth-neutral robustness checks that distinguish motive from a mechanical donor-breadth artifact.

**Committee-type classification.** DIME `contributor_type = "C"` is the full set of registered committee filers, not a corporate subset; its only sibling value is `"I"` (individuals). To recover each committee's FEC type I map the DIME committee identifier (`bonica_cid`) to its modal FEC committee identifier (`efec_comid`) and join the FEC committee master, reading committee type (`CMTE_TP`), organization type (`ORG_TP`), and designation (`CMTE_DSGN`). Because `efec_comid` is populated only from 2012 onward while `bonica_cid` is stable across the panel, a committee's 2012-or-later FEC identifier backfills the classification of its earlier-cycle contributions. Coverage is therefore era-dependent: 86 percent of committee-race observations are classified in 2012–2020 but only 22 percent before 2002, and roughly half of all close-bandwidth observations (overwhelmingly committees that ceased filing before 2012) remain unclassified. The `CMTE_TP` field (which distinguishes candidate and party committees from PACs) is well populated where present; `ORG_TP` (which marks a connected committee's sponsor type as corporate, labor, trade, membership, or cooperative) is frequently blank, so a committee positively coded corporate by `ORG_TP` is a strict lower bound on business committees. I recover a broader business-connected set from the presence and text of the connected-organization name (`CONNECTED_ORG_NM`). Two facts about the `CMTE_TP` N/Q category

Table P1: Composition of the committee sample within the analysis bandwidth.

Share of committee-race observations ($|rv| < 14.3$ pp, the pre-extension headline bandwidth this donor-type analysis retains) by FEC committee type.

| Committee type | Share of close-bandwidth observations |
|---|---|
| Candidate committees (CMTE_TP H/S/P) | 23.7% |
| Party committees (CMTE_TP X/Y/Z) | 21.0% |
| N/Q-typed PACs (88% carry a connected-org name) | 3.9% |
| Leadership PACs (CMTE_DSGN = D) | 1.7% |
| Trade (ORG_TP = T) | 1.5% |
| Membership / labor / cooperative (ORG_TP M/L/V) | 0.9% |
| Corporate (strict, ORG_TP = C/W) | 0.7% |
| Unmatched or unclassified (mostly pre-2012) | 46.6% |

govern everything below. First, N and Q denote *nonqualified* and *qualified* PAC filing status (a registration-age and breadth-of-support distinction), not the absence of a sponsor: connected committees file as N/Q too, and sponsorship is recorded only in ORG_TP and CONNECTED_ORG_NM. Second, applying that correction, 2,797 of the 3,105 committees that an N/Q-based rule would label "non-connected" carry a connected-organization name and belong to the business-connected set; the genuinely non-connected residual is 308 committees. Table P2 reports the naive N/Q bucket for completeness alongside the corrected, mutually exclusive rows; the body's Table 4 carries the corrected rows.

Party and candidate committees together account for roughly two-fifths of all close-bandwidth observations; positively coded corporate committees are under one percent, which is what makes a "corporate PAC" label for this sample a misnomer.

**Redirection by committee type.** Each row of Table P2 is the body's headline specification (rdrobust, triangular kernel, local-linear, MSE-optimal bandwidth, race-clustered standard errors) run on the indicated committee-type subset. The pooled re-run on the classified sample reproduces the body headline exactly (+0.375), validating the classification pipeline. Every committee type redirects toward the winner, from +0.25 (strict corporate) to +0.46 (non-connected). Two rows are diagnostic and overlap the others: the N/Q-typed row pools every committee filing under CMTE_TP N/Q, 88 percent of whose observations belong to committees with a connected-organization name that the recov-

ered classification assigns to the business-connected row; and the strict-corporate row is the noisy 657-committee `ORG_TP` = C slice that understates business giving. The mutually exclusive recovered rows (business-connected, labor, and non-connected) are the ones the body's Table 4 reports. Excluding party and candidate committees (the "PAC-only" row) leaves the redirection essentially unchanged at +0.357 on the classified subsample; the Appendix Table B3 variant, which retains committees whose FEC type is unclassified, is +0.377; the sample difference, not instability, separates the two. Genuine type differences appear on partisan reorientation: business-connected types reorient toward the winning party ($\approx 0.06$), whereas labor reorients negligibly (0.023, about one and a half standard errors from zero) and the non-connected residual reorients at 0.039 (standard error 0.015), between the labor and business values.

Every purpose class redirects strongly, and every class converts: ideological and single-issue committees fund the opponent who beat their pick at +0.081 and +0.044, smaller than business's +0.134 but unambiguously positive. The CRP business class is narrower than the FEC business-connected set and is estimated on the bridged subsample, so its +0.258 is not comparable to the FEC-classified +0.374; comparisons in the text stay within one classification. Where the classes differ on conversion, the difference is baseline portfolio structure rather than professed purpose. Stratifying committees on a predetermined two-party-giving measure (the minority-party share of the committee's non-focal cycle-t party-linked dollars) puts almost all conversion inside the two-party stratum: among committees whose baseline giving already spans both parties, conversion is statistically indistinguishable across classes (business +0.186, ideological +0.185, single-issue +0.179; pairwise differences under one standard error), while among committees giving to one party only it is small (at most +0.028, for business, and not distinguishable from zero in the other classes). The stratifier is smooth at the cutoff (twelve balance discontinuities on stratum membership, the largest 0.024, none significant at five percent), so the decomposition is not an artifact of treatment-side sorting. That

concentration says who pays the crossover premium; it does not make the premium the routine allocation of a two-party portfolio. Within the two-party stratum the opponent's funding jumps at the cutoff, and the opponent's identity, party, and pre-vote proximity are the same on both sides of it, so these committees fund the opponent because the opponent won. What the one-party strata show is Section 2's proximity term ($V_I \rho_c$) operating on the crossover margin: committees whose baseline giving stays on one side of the aisle rarely cross it even when the opponent wins, the bounded ideology channel the body reports alongside the office premium. Two limits bound the reading. Portfolio structure is a committee's own choice, so these are conditional-on-choice comparisons: ideology may operate precisely by keeping a committee's giving on one side of the aisle (the party-lockedness Fouirnaies and Hall (2014) describe for labor unions), and the within-stratum estimates are heterogeneity-tier evidence, not main specifications. What the classes differ on unconditionally is retention: ideological and single-issue committees fund their own winners at higher rates than business committees in every stratum.

**Individual-donor panel.** I build the individual panel from DIME `contributor_type = "I"` using the identical race definitions and dominant-amount pick rule as the committee panel, so that only the donor universe differs. The individual panel keeps earmarked contributions (transaction type 15E), which the committee panel excludes (Appendix Section A.1): on the individual side an earmarked gift routed through ActBlue or WinRed is the donor's own contribution to the candidate, the t+2 choice the design measures, whereas on the committee side the same code marks a conduit transmitting another person's money; individuals file no independent expenditures, so no other filter applies. The races are extracted from the committee analytic panel itself, making the race set byte-identical, and the window is the committee panel's: focal cycles 1990 through 2022, with t+2 outcomes through 2024. Because the individual data are large (the 2020 cycle alone carries roughly 22 million itemized House contributions), I assemble the panel in a disk-backed columnar database, aggregating per donor-race, assigning the pick,

and joining the t+2 outcomes, and collect only the matched donor-race observations: 20,365,460 across 6.28 million individual donors. Three-fifths of the observations (61.3 percent) come from the 2020 and 2022 focal cycles, the conduit-era small-donor flood that any analysis of this population must confront. The committee benchmarks re-derived in the same scripts reproduce the body's canonical estimates exactly, validating the shared pipeline.

**Size strata.** Donors are stratified by giving to the focal candidates at fixed dollar cutoffs: small under $200, mid $200–500, large over $500. The cutoffs are fixed rather than sample-relative because the era composition would otherwise move the strata themselves (four-fifths of 2020s donor-race observations come from donors giving under $100, so within-sample terciles would collapse below the itemization threshold), and the $200 boundary coincides with the FEC itemization threshold, which governs interpretation below it.

**Itemization censoring.** FEC rules mandate itemization only for donors whose contributions to a recipient exceed $200 in a cycle; below that threshold a contribution appears in DIME only when the committee itemizes voluntarily or the gift passes through a conduit (ActBlue, WinRed) that itemizes everything. Small-donor rows are therefore a selected subset of small giving, and a small donor's *total* observed giving is censored in a way a large donor's is not (every contribution by a donor above the threshold is itemized by law). Two consequences follow for the analysis: the total-giving outcome is interpreted only on the large stratum (and secondarily mid), where observation is complete; and the small-donor redirection and reorientation estimates are read as statements about observed small giving, with the breadth-neutral reorientation share (a within-portfolio fraction unaffected by how much of the portfolio is observed at the margin) carrying the primary weight.

The small-donor reorientation is precisely estimated (0.005, standard error 0.002) but an order of magnitude below the committee level, and its raw sign varies across focal

subwindows ($-0.002$ on 1990–2018, +0.005 on the full window), both values substantively zero. Only the mid and large budget coefficients are distinguishable from zero ($p = 0.03$ and $p = 0.003$); across the twelve stratum-by-outcome tests of this table, the large coefficient survives a Bonferroni correction ($0.003 \times 12 = 0.03$) and the mid coefficient does not ($0.029 \times 12 = 0.35$), so only the large-stratum budget response survives the correction.

**Budget margins.** The large-stratum budget response decomposes into both margins: a narrow win raises the probability of any House giving at t+2 by +0.054 (standard error 0.022) and raises totals among donors active in both cycles by +0.112 log points (standard error 0.046; post-treatment conditioning caveat as for the conditional-on-active check below). The mid stratum shows the same signs at lower precision (+0.049 and +0.038, neither distinguishable from zero). The committee panel is null on both margins (+0.011 and $-0.019$, neither distinguishable from zero), consistent with Appendix Section Q's exit-robust battery: organizational budgets are reallocated, individual money is mobilized and demobilized. The large-stratum total effect is bandwidth-stable (+0.41 at half the MSE bandwidth and +0.54 at double, both distinguishable from zero at the five percent level).

**The ActBlue-era read.** Restricting to focal cycles 2020–2022 (the small-donor-flood era), the gradient survives and sharpens: winner redirection runs small +0.145, mid +0.339, large +0.485 (the committee 2016–2024 era estimate is +0.377), the large-donor budget response persists at +0.425 (about 1.9 standard errors from zero), and the small-donor reorientation remains substantively zero (0.007). The relocation claim is not an artifact of the pre-conduit era.

**Bandwidth stability by stratum.** Halving or doubling each stratum's MSE-optimal bandwidth (h/2, 3h/4, 3h/2, 2h) leaves the gradient intact: winner redirection stays within five percent of its baseline in every stratum (small +0.140 to +0.147, mid +0.285 to +0.296, large +0.463 to +0.471), the reorientation share keeps the break between small donors and everyone else (small 0.003 to 0.004, mid and large 0.040 to 0.047), and the

large-stratum budget response runs +0.41 to +0.54 (each distinguishable from zero at the five percent level) while the mid-stratum response, +0.26 to +0.36, is distinguishable from zero only at the wider bandwidths and the small-stratum response stays within $\pm 0.07$ of zero throughout.

**Breadth-neutral robustness and the pre-period trend.** The binary `gave_to_winner_t2` gradient could in principle be a mechanical artifact: donors who give to more candidates are more likely to give to any one of them, the winner included. Three checks establish that the gradient reflects motive rather than donor breadth. First, the partisan-reorientation outcome is a share of the donor's portfolio and is therefore invariant to how many candidates the donor supports; on it the gradient is small 0.005, mid 0.042, large 0.042 (committee 0.062): small donors barely realign, while mid-sized and large donors realign at two-thirds of the committee level. Second, small donors are not less active at t+2 than mid-sized donors (activity rates 57–59 percent against 46–53), so the weak redirection of small donors is not attrition or churn. Third, conditioning on donors active at t+2 (which removes the breadth/retention channel at the cost of post-treatment conditioning) leaves the gradient intact and sharper (+0.246, +0.559, +0.648 across the strata). Like its committee counterpart (Appendix Section J), the reorientation share fails the pre-period balance test, and for the same reason: the two sides of the cutoff are different donors, leaning toward opposite parties, whose own-party sorting was already diverging. For donors observed two cycles earlier (39 percent of donor-race rows), the pre-period trend is +0.023 pooled, and roughly two-thirds of it is the donor's own focal-race giving inside the share: removing the focal candidates at both endpoints leaves +0.008, and the window ending two cycles before the race gives +0.010. The focal-excluded trend rises with donation size (+0.005, +0.012, +0.018 across the strata), so netting each donor's drift out steepens the gradient rather than flattening it: the trend-adjusted values (computed on the focal-inclusive convention, matching the post-period share) are 0.005, 0.073, and 0.086, strictly monotone, with the large-donor value

equal to the committee panel's focal-inclusive 0.087 (the small-donor value changes sign under adjustment and stays an order of magnitude below the rest). The breadth-neutral share, raw or trend-adjusted, is the primary evidence; the conditional-on-active estimate corroborates it.

**Office premiums by stratum.** Applying the Section 4.1 reparametrization (`gave_to_pick_t2`, `gave_to_opp_t2`; crossover reported loss-side over win-side) to the individual panel separates retention from conversion. The committee values re-derived in the same pass reproduce the body exactly (+0.200 own-pick, +0.069 crossover).

Large donors match the committee own-pick premium almost exactly, and the crossover premium is size-monotone and precisely estimated for mid and large donors, but at +0.011 it is a sixth of the committee value, and it is exactly zero for small donors. The comparison also explains why the large-donor *composite* (+0.472) exceeds the committee composite (+0.375): the composite subtracts the loss-side crossover level from the win-side continuation level, so the class that crosses over least mechanically posts the largest composite. Reading the premiums rather than the composite, individual winner-following is dominated by retention of successful picks; outright conversion to a former opponent (the access component ideology cannot generate) remains overwhelmingly organizational. One robustness caveat attaches to the individual own-pick premiums: dropping the 2020 focal cycle, which carries two-fifths of the individual panel's kernel weight, cuts the small-stratum premium by more than half (to +0.026, not distinguishable from zero) and the pooled individual premium by 16 percent, while the mid and large strata and every winner-following estimate are stable to the same drop.

**Who the graded donors are.** For the 2020–2024 cycles (the window covered by the locally held raw itemized files), I profile close-race donors ($|\text{rv}| < 15$) by the occupation string carrying each donor's largest total House giving. Policy-exposed occupational categories rise steeply across the strata: business executives and owners 6.3, 10.5, and 15.2 percent (small, mid, and large), attorneys 2.3, 6.1, and 7.5, finance and real estate 1.2, 2.6,

and 4.0, physicians 1.4, 2.9, and 3.5 (jointly roughly a third of large donors against an eighth of small donors), while educators and engineers are flat or declining shares. Two reporting caveats: the "not employed" category (28 percent of small donors and 22 percent of mid and large donors) is inflated by conduit reporting conventions that record blank employers as not employed, and retirees' prior occupations are unobserved; both work against finding an occupational gradient, so the recovered gradient is, if anything, understated. The profile is descriptive and era-scoped, but it matches the behavioral evidence: the donors whose giving follows winners, respond on the budget margin, and carry the trace crossover are disproportionately the professionals and principals with direct policy exposure (the "investors" of the classic donor-motive typology, Francia et al., 2003), where the small-donor pool is dominated by donors with no occupational stake recorded.

**Density composition.** The observation-weighted McCrary test rejects on the individual panel as it does on the committee panel (Appendix Section J, Table J1), with an instructive sign contrast: the individual-panel density is higher on the *losing* side (t-statistic $-33.1$ on the full window; $-13.5$ among donors giving under \$200 and $-41.6$ among donors giving over \$500), against a committee panel denser on the winning side (+9.2), because individual money concentrates on close-race challengers who narrowly lose, where committee money concentrates on the incumbents who narrowly beat them. As in Section J, these are mechanical composition facts about who funds whom in close races, not evidence of running-variable manipulation: donors do not determine election outcomes, the race set is byte-identical to the committee panel's, and the appropriate validity evidence is covariate balance (de la Cuesta and Imai, 2016).

Scripts: `code/33_pac_type_heterogeneity.R`, `code/34_individual_donor_rdd.R`, `code/35_donor_gradient_validation.R`, `code/54_individual_battery.R`–`code/61_donor_gradient_validation_ext.R`; classification audit, CRP external classification, and conversion-composition analyses in `code/62_nq_ideological_anatomy.R`–

`code/64b_crp_loser_estimates.R`.

# Q Total-giving null: the exit-robust budget battery

The body's third co-primary outcome, the change in log total House giving, is a precise zero at the cutoff (+0.032 log-units, standard error 0.131; Table 3, column 3), but its 95 percent confidence interval is wide: the upper limit, +0.273 log-units, corresponds to $\exp(0.273) - 1 \approx 31$ percent expansion, wide enough that a budget-expansion hybrid could in principle inhabit it. This appendix shows that the width is an artifact of the logarithm amplifying committee exit rather than genuine imprecision about the quantity the hybrid predicts, and that on every margin matching the hybrid's prediction the budget is fixed to within 5.1 to 7.3 percent.

Roughly 14 percent of committees make no House contribution at all by t+2 (the active-at-t+2 rate is 0.863). For each such committee the log change is $\log(1) - \log(\text{total}_t + 1) = -\log(\text{total}_t + 1)$, a large negative number; the full-sample delta-log distribution has mean $-1.39$ and standard deviation 3.39, and this heavy left tail (committees disengaging entirely) is what inflates the confidence interval. Exit is the opposite of the expansion a hybrid requires, so the wide unconditional interval does not keep the hybrid alive. Three exit-robust reformulations isolate the expansion question directly.

The log-change family (rows A–F) holds the expansion bound in the 27–35 percent range throughout, including under the two analysis-of-covariance specifications. The analysis-of-covariance result is diagnostic: controlling for the committee's baseline giving level barely moves the standard error (0.131 to 0.134), which shows the residual variance is not cross-committee level heterogeneity (which the baseline absorbs) but the within-committee exit channel (which it does not). Winsorizing at the first and ninety-ninth percentiles does not tighten the bound either, because the exit mass is roughly 14 percent of committees, not one percent.

The exit-robust family (rows E, G, H) collapses the bound to 5.1 to 7.3 percent. Conditioning on committees that remain active at t+2 gives a loss effect of +1.9 percent (expansion ruled out above 7.3 percent). A within-committee fixed-effects specification, comparing a committee's own close wins to its own close losses and absorbing each committee's exit propensity (adjusted $R^2 = 0.53$), gives +0.6 percent (expansion ruled out above 5.5 percent). Expressing the change as a share of the committee's prior budget, which caps an exit at $-100$ percent rather than sending it to $-\infty$, gives +0.1 percent (expansion ruled out above 5.1 percent). All three are precisely estimated near-zeros.

The within-committee specification is identified off the 19,326 committees (28.6 percent of the close-bandwidth sample) that appear in close races on both the winning and losing sides, the same multi-race sub-population that identifies the within-PAC fixed-effects sensitivity (Appendix Section M), and so inherits its sub-population caveat; the active-only specification conditions on a post-treatment survival indicator, defensible because differential attrition at the cutoff is small and roughly balanced (Section J) but post-treatment conditioning nonetheless. Neither is decisive alone. Their convergence with the share-of-budget margin (which is full-sample and conditions on nothing post-treatment) is what licenses the conclusion. Three reformulations that remove the exit tail in three different ways all return a change within roughly two percent of zero. The redirection toward the winner is financed by substitution within that near-fixed budget, the budget-constrained allocation of the model in Section 2. The bound is on the average loss effect on total giving; it does not bound heterogeneity, so a minority of committees could expand while others contract. Script: `code/31_total_giving_null.R`.

# R The lobbying channel: directed-lobbying redirection on the public-firm subsample

The body reports that the redirection toward the winner appears in a second, behaviorally independent channel, directed lobbying, on the subsample where lobbying can be observed. This section gives the construction and the full battery.

Lobbying expenditures are observable, and matchable to specific legislative targets, only for committees sponsored by publicly traded firms. I build a firm-by-focal-race panel that mirrors the headline committee design exactly: the unit is a sponsoring firm (Compustat `gvkey`) crossed with a contested two-candidate House focal race; the running variable is the signed vote-share margin from the firm's perspective (positive if the firm's supported candidate won the focal race, negative if that candidate lost), with the cutoff at zero and a dominant-amount pick rule for firms that gave to both top-two candidates. Prior support is matched on the Federal Election Commission candidate identifier (`CAND_ID`), observed for 100 percent of both winners and losers in close focal races; matching on a congressional-service identifier instead would restrict the loser-backer arm to former incumbents (30 percent coverage) and induce differential selection at the cutoff. The outcome is the firm's directed lobbying of the focal winner over the t+2 cycle window, summed from the firm-by-legislator-by-year lobbying panel of Kim et al. (2026), zero-filled when the firm did not lobby the target (the same zero-treated-as-zero convention used for the headline outcomes). Restricting to focal cycles 2008 through 2016 (so the t+2 outcome window closes by 2018, the panel's last year) yields 116,718 firm-by-race observations across 1,074 firms and 1,669 races; the near-cutoff sample at the primary bandwidth is ample (effective N $\approx$ 26,500). Because the binary "did the firm lobby the winner at all" indicator is saturated (roughly 69 percent), the primary outcome is lobbying *intensity*, the log of one plus directed lobbying dollars.

The primary directed-lobbying discontinuity is positive, significant, and stable across

fixed bandwidths: +1.16 (h = 5), +1.24 (h = 10), +1.23 (h = 14.3), and +1.17 (h = 20), with precision improving as the effective sample grows. The two giving cross-checks reproduce the redirection inside the same donation data (+0.17 binary, +1.41 log dollars), confirming the subsample behaves like the headline; the binary cross-check, +0.171, is about half the headline +0.375, as expected for a different sample (public-firm committees, 2008–2016) and a different giving measure, and I do not read the two magnitudes as equal. The placebo (directed lobbying of the focal loser, who is out of office at t+2) is a null ($-0.295$, standard error 0.433), so the discontinuity is specific to the candidate who holds the seat.

Dropping the off-year leaves the primary essentially unchanged (+1.16), so it is not a window-pooling artifact. Conditioning on firms that lobby at all isolates the intensive margin and if anything strengthens the discontinuity (+1.42), while the giving cross-checks reproduce and the placebo stays null on the active subsample. Directed lobbying of the winner's committees also jumps (+0.15); at base rate 0.64 this directed-and-windowed measure is no longer saturated (unlike the firm-level any-committee-connection indicator, which is saturated at 80 to 94 percent across sitting legislators and is therefore unusable as an RDD outcome), and it corroborates the intensity result.

The firm-side predetermined covariates (prior giving and split-giving behavior) are balanced at the cutoff. The incumbent-reelected jump is not a firm-side confound but the same mechanical winner-backing asymmetry that drives the McCrary rejection in the headline panel: firms back winning incumbents more than losing challengers, so incumbent-reelected races skew to the winner side of the cutoff, and the headline already shows the redirection holds within each incumbency stratum (Appendix Section B.2). The density result is the cleanest demonstration of that mechanism. The firm-replicated panel rejects density continuity (test statistic 6.74), mirroring the headline panel's rejection, because each race contributes one observation per firm-backer and winners attract more firm-backers. But when each race contributes its winner-side margin and its loser-side margin once each (removing the firm-count replication), the underlying race-level

margin density is continuous (test statistic 0.55), so the rejection is purely the firm-count replication artifact, a cleaner illustration of the de la Cuesta and Imai (2016) reading than the headline panel provides.

The subsample is 2008–2016 focal cycles on committees sponsored by publicly traded firms; the test is mechanism confirmation on the subset where both giving and lobbying are observable, not an extension of the headline universe, and the lobbying magnitude is not the giving magnitude. Scripts: `code/27_build_lobbying_bridge.R`, `code/28_lobbying_rdd.R`, `code/29_lobbying_rdd_robustness.R`.

# S Replication materials

Replication materials (the analytic panel, all numerical outputs underlying the body and appendix tables, and the canonical-script logs that produced them) are available from the author on request.

Table P2: Winner redirection, loser-loyalty drop, and partisan reorientation by committee type.

`gave_to_winner_t2`, `gave_to_loser_t2`, and the sign-aligned partisan-reorientation outcome; magnitudes in the body's $|\hat{\beta}|$ convention. Rows above the rule are mutually exclusive under the recovered classification; rows below it are diagnostic slices that overlap them.

| Committee type | →winner | →loser | \|reorient\| |
|---|---|---|---|
| Party committees | +0.401 | −0.217 | 0.038 |
| Leadership PACs | +0.355 | −0.211 | 0.079 |
| Candidate committees | +0.347 | −0.187 | 0.072 |
| Labor | +0.367 | −0.196 | 0.023 (n.s.) |
| Membership | +0.311 | −0.145 | 0.089 |
| Trade | +0.285 | −0.167 | 0.113 |
| Business-connected (recovered, 3,453 committees) | +0.374 | −0.213 | 0.065 |
| Non-connected (no sponsor, 308 committees) | +0.456 | −0.192 | 0.039 |
| Corporate (strict `ORG_TP` = C, 657 committees) | +0.253 | −0.146 | 0.134 |
| N/Q-typed (all `CMTE_TP` N/Q; 88% sponsored) | +0.404 | −0.220 | 0.050 |
| PAC-only (excludes party + candidate committees) | +0.357 | −0.199 | 0.074 |
| Pooled (= body headline) | +0.375 | −0.227 | 0.062 |

**External purpose classification (CRP).** The FEC fields above classify structure (who sponsors a committee), not purpose. The federal-level proxy standard since Snyder (1990), used at the federal level by Fouirnaies and Hall (2014), treats individuals and non-connected PACs as consumption-motivated contributors and every other interest group as investor-motivated; the sponsor-name audit above shows why an N/Q implementation of that proxy is fragile. As a direct measure of purpose I therefore adopt the Center for Responsive Politics committee classification, which hand-codes every FEC committee: `RecipCode` assigns business (PB/OB), labor (PL/OL), and ideological (PI/OI) types, and the industry catcode's sector J marks Ideology/Single-Issue committees. The strictest definition below (sector J excluding the partisan-electoral and leadership families, J11–J24) recovers the canonical single-issue set (the National Rifle Association, NARAL, the National Organization for Women, the League of Conservation Voters). Classification is joined committee-by-cycle, since an FEC identifier can change identity across cycles; it covers 47.8 percent of panel observations through the 2012+ FEC-identifier bridge, and the pooled headline re-estimated on the bridged subsample reproduces the panel value exactly (+0.374), with the two premiums within one standard error of their full-panel values (continuation +0.226 against +0.200, crossover +0.062 against +0.069; Table P2b), licensing the comparison.

**Table P2b: Redirection and office premiums by CRP purpose class (bridged subsample).** Body specification throughout; premiums per the conventions of body Section 4.1 (crossover reported loss-side over win-side); reorientation in the body's $|\hat{\beta}|$ convention. Committee counts: business 1,824, labor 332, ideological 3,723, single-issue 531. Business and labor are disjoint classes; ideological (RecipCode PI/OI) and single-issue (sector J excluding the J11–J24 partisan-electoral and leadership families) are nested definitions of one purpose concept, from more to less inclusive, not disjoint classes.

| CRP class | →winner | →loser | Continuation | Crossover | \|reorient\| |
|---|---|---|---|---|---|
| Business | +0.258 (0.036) | −0.164 (0.040) | +0.231 (0.051) | +0.134 (0.028) | 0.119 (0.016) |
| Labor | +0.499 (0.042) | −0.315 (0.050) | +0.259 (0.063) | +0.077 (0.020) | 0.005 (0.006, n.s.) |
| Ideological | +0.451 (0.022) | −0.231 (0.039) | +0.265 (0.047) | +0.081 (0.012) | 0.042 (0.008) |
| Single-issue | +0.494 (0.028) | −0.237 (0.041) | +0.276 (0.050) | +0.044 (0.015) | 0.021 (0.008) |
| Bridged pooled | +0.374 (0.019) | −0.195 (0.026) | +0.226 (0.036) | +0.062 (0.011) | 0.054 (0.007) |

Table P3: The four-outcome battery for individuals versus committees (focal 1990–2022).

Body specification throughout (rdrobust, triangular kernel, local-linear, MSE-optimal bandwidth, race-clustered SEs); magnitudes in the body's $|\hat{\beta}|$ convention for the reorientation column.

| Donor group | →winner | →loser | \| reorient \| | Δ log total giving |
|---|---|---|---|---|
| Individuals (all) | +0.201 (0.013) | −0.123 (0.023) | 0.011 (0.003) | +0.14 (0.15) |
| Individuals, small (< $200) | +0.144 (0.014) | −0.089 (0.025) | 0.005 (0.002) | +0.04 (0.15) |
| Individuals, mid ($200–500) | +0.300 (0.015) | −0.157 (0.018) | 0.042 (0.004) | +0.33 (0.14) |
| Individuals, large (> $500) | +0.472 (0.014) | −0.255 (0.026) | 0.042 (0.004) | +0.48 (0.15) |
| Committees (full window) | +0.375 (0.016) | −0.227 (0.023) | 0.062 (0.008) | +0.03 (0.13) |

Table P4: Own-pick and crossover premiums, individuals versus committees.

| Donor group | Own-pick premium | Crossover premium |
|---|---|---|
| Individuals (all) | +0.085 (0.026) | +0.002 (0.001, p = 0.06) |
| Individuals, small (< $200) | +0.058 (0.027, p = 0.05) | −0.000 (0.001, p = 0.59) |
| Individuals, mid ($200–500) | +0.147 (0.024) | +0.006 (0.001) |
| Individuals, large (> $500) | +0.214 (0.033) | +0.011 (0.002) |
| Committees | +0.200 (0.031) | +0.069 (0.010) |

Table Q1: Total-giving null across eight specifications.

Rows A–F keep the committee-exit mass in the outcome and vary the transform (delta log, inverse-hyperbolic-sine change, two analysis-of-covariance specifications controlling for the baseline giving level, and a winsorized delta log); rows E, G, H remove or bound the exit mass (conditioning on committees active at t+2, within-committee fixed effects, and the share-of-budget percentage change winsorized at p1/p99). Specifications A–F, E, H are sharp `rdrobust` (Calonico et al., 2014; triangular kernel, MSE-optimal bandwidth, race-clustered HC1, robust bias-corrected CI); G is a within-committee LPM (`fixest::feols`, committee and cycle fixed effects) with the running variable interacted with the loss indicator. "Expansion bound" is the upper 95 percent confidence limit (the side the budget-expansion hybrid predicts) translated to a percentage via $\exp(\cdot) - 1$ for the log/IHS rows and directly for the share row.

| Spec | Formulation | Point (SE) | $p$ | 95% CI | Expansion bound | $N_h$ |
|---|---|---|---|---|---|---|
| A | $\Delta \log$ total giving (baseline) | +0.032 (0.131) | 0.90 | [−0.240, +0.273] | 31.4% | 735,469 |
| B | $\Delta \operatorname{asinh}$ total giving | +0.044 (0.138) | 0.82 | [−0.240, +0.301] | 35.2% | 748,268 |
| C | ANCOVA $\log(t+2)$ on baseline $\log$ | −0.013 (0.134) | 0.73 | [−0.309, +0.217] | 26.6% | 640,168 |
| D | ANCOVA $\operatorname{asinh}(t+2)$ on baseline | −0.006 (0.141) | 0.78 | [−0.316, +0.238] | 27.1% | 647,447 |
| F | $\Delta \log$ winsorized (p1/p99) | +0.034 (0.130) | 0.88 | [−0.235, +0.274] | 31.5% | 737,754 |
| E | $\Delta \log \mid$ active at t+2 | −0.019 (0.027) | 0.37 | [−0.076, +0.028] | 7.3% | 506,003 |
| G | within-committee FE, loss coefficient | +0.006 (0.024) | 0.80 | [−0.041, +0.054] | 5.5% | 735,469 |
| H | share-of-budget % change, winsorized | −0.001 (0.025) | 0.93 | [−0.051, +0.046] | 5.1% | 816,183 |

Table R1: Directed-lobbying RDD on the public-firm subsample (full panel).

Sharp `rdrobust` (Calonico et al., 2014; triangular kernel, MSE-optimal bandwidth, race-clustered HC1 standard errors, robust bias-corrected inference), focal cycles 2008–2016. All outcomes are measured toward the focal winner at t+2 except the placebo (toward the out-of-office loser). The two cross-checks re-estimate the giving redirection inside the same firm-by-legislator donation data.

| Outcome (toward winner, t+2) | Point | Robust SE | $p$ | 95% CI | $h$ |
|---|---|---|---|---|---|
| log directed lobbying $ (PRIMARY) | +1.245 | 0.445 | 0.003 | [+0.46, +2.20] | 12.3 |
| lobbied a bill the winner sponsored (7.6%) | +0.020 | 0.012 | 0.064 | [−0.001, +0.044] | 15.8 |
| any lobbying of winner (saturated, 69%) | +0.117 | 0.046 | 0.006 | [+0.04, +0.22] | 11.8 |
| log lobbying, winner − loser (net) | +1.549 | 0.817 | 0.030 | [+0.17, +3.37] | 12.0 |
| gave to winner (cross-check, binary) | +0.171 | 0.031 | $7 \times 10^{-9}$ | [+0.12, +0.24] | 12.9 |
| log donations to winner (cross-check) | +1.410 | 0.253 | $7 \times 10^{-9}$ | [+0.97, +1.96] | 13.1 |
| log lobbying of loser (PLACEBO) | −0.295 | 0.433 | 0.37 | [−1.24, +0.46] | 14.3 |

Table R2: Robustness of the primary directed-lobbying discontinuity.

Same estimator. Block (i) restricts the outcome to the single post-election year; block (ii) restricts the firm sample to the 752 firms with any positive lobbying in the panel (conditioning out the firm-level extensive margin); row (iii) is directed lobbying of a committee on which the winner sits.

| Specification | Point | Robust SE | $p$ | $h$ |
|---|---|---|---|---|
| (i) log lobbying of winner, t+2 year only [full] | +1.164 | 0.420 | 0.003 | 11.8 |
| (i) log lobbying of winner, t+2 year only [active firms] | +1.332 | 0.484 | 0.003 | 11.6 |
| (ii) log lobbying of winner [active firms, PRIMARY] | +1.424 | 0.514 | 0.003 | 12.0 |
| (ii) lobbied a bill winner sponsored [active firms] | +0.024 | 0.014 | 0.064 | 15.7 |
| (ii) gave to winner [active firms, cross-check] | +0.181 | 0.030 | $2 \times 10^{-10}$ | 13.5 |
| (ii) log donations to winner [active firms, cross-check] | +1.491 | 0.243 | $2 \times 10^{-10}$ | 13.6 |
| (ii) log lobbying of loser [active firms, PLACEBO] | $-0.356$ | 0.520 | 0.36 | 14.0 |
| (iii) directed committee lobbying of winner | +0.148 | 0.048 | 0.001 | 12.1 |

Table R3: Running-variable validity on the firm-by-race panel.

Firm-side predetermined covariates run as `rdrobust` outcomes (balanced at $|t| < 2$); density via `rddensity` at the cutoff.

| Check | Estimate / statistic | Balanced / continuous |
|---|---|---|
| balance: log firm prior giving at t | $-0.055$ ($\lvert t\rvert = 0.71$) | yes |
| balance: gave to both top-two (dual pick) | $-0.007$ ($\lvert t\rvert = 0.54$) | yes |
| balance: open seat | $-0.007$ ($\lvert t\rvert = 0.11$) | yes |
| balance: incumbent reelected | +0.659 ($\lvert t\rvert = 9.19$) | no (mechanical) |
| density: firm-by-race panel | $T = 6.74$ ($p = 2 \times 10^{-11}$) | rejects |
| density: race-by-side deduplicated | $T = 0.55$ ($p = 0.585$) | continuous |