

The Luck of the Draw: New Evidence on Windfalls from Thousands of Randomized Income Shocks

Torun Dewan* Daniel Rubenson†

March 14, 2026

Abstract

Does money improve organizational performance, and under what conditions? We exploit the random draw of the English FA Cup—the oldest knockout competition in football (soccer)—which assigns financial windfalls to lower-division clubs drawn to play at the stadiums of much larger opponents. Using over 5,000 club-season observations spanning 1952 to 2022, we find that clubs receiving these windfalls climb approximately 18 positions through the English league pyramid over ten years, are less likely to be relegated (demoted to a lower division), and more likely to be promoted. The effect is present even among clubs that lose the match—receiving money without any confounding boost from an upset victory—providing strong evidence that the mechanism is financial rather than psychological. Consistent with this interpretation, the treatment effect across decades tracks the changing financial salience of gate receipts, rising as clubs became more financially constrained and declining as television revenues transformed the economics of English football. The effect on upward mobility is moderated by organizational capacity: clubs with greater pre-existing infrastructure benefit more from windfalls. This moderation applies to upward movement but not to protection against relegation, suggesting that a financial buffer alone can prevent decline, while turning resources into advancement requires the capacity to spend them well.

Keywords: financial resources; organizational performance; organizational capacity; natural experiment; windfalls; FA Cup

*Department of Government, London School of Economics. Email: t.dewan@lse.ac.uk

†Department of Political Science, University of Toronto. Email: daniel.rubenson@utoronto.ca

1 Introduction

Does money improve organizational performance? The question is fundamental to management, politics, economics, and organizational theory. Firms invest in talent, capital expenditure, and strategic initiatives; political parties fund campaigns and mobilize supporters; universities recruit faculty and build facilities. More financial resources should, in principle, translate into better performance. Yet establishing this relationship causally has proven difficult. Organizations that perform well attract more resources, creating a feedback loop that confounds naïve estimates of the effect of money on outcomes. Unobserved factors—managerial quality, institutional culture, strategic vision—simultaneously drive both resource acquisition and performance.

This identification challenge pervades several important literatures. In management, the resource-based view of the firm (Barney 1991) identifies resources as a source of competitive advantage, but the endogeneity of resources and performance makes it difficult to test whether exogenous increases in financial capital actually improve outcomes. The organizational slack literature debates whether excess resources enable productive experimentation or breed complacency (Cyert and March 1963, Nohria and Gulati 1996), but lacks truly exogenous variation in slack. In corporate finance, research on cash windfalls suggests that unexpected money can lead to agency problems rather than improved performance (Blanchard et al. 1994, Jensen 1986). In political science, a large literature debates whether campaign spending actually causes electoral success (Jacobson 1978, Gerber 1998, Levitt 1994). And in development economics, the “resource curse” literature finds that natural resource wealth often fails to improve economic or political outcomes (Sachs and Warner 1995, Ross 2001). Across these literatures, the core question is the same—does money help?—but credible causal evidence remains scarce.

In this paper, we exploit a natural experiment that addresses many of these challenges: the random draw of the Football Association (FA) Cup, the oldest and most prestigious knockout football (soccer) competition in the world. Each year since 1871, clubs from across the English football league system—from billion-dollar Premier League franchises to tiny non-league clubs—are entered into a single-elimination tournament. In the third round proper, when the top-division clubs enter, matchups are determined by a random draw. When a small club is drawn to play away at a much larger club, it receives a share of the gate receipts from a match played in a large stadium—a financial windfall that can transform the smaller club’s finances. This windfall is assigned by a random draw that is orthogonal to the smaller club’s quality, recent performance, or any other confounding factor.

This setting offers several advantages for studying the causal effect of financial resources on organizational performance. First, treatment is randomly assigned. Second, the treatment—money from gate receipts—is clearly defined and easily measured. Third, the outcomes—movement through the league pyramid, relegation, and promotion—are unambiguous and objectively measured. Fourth, the organizations under study (football clubs) are relatively homogeneous in their objectives: they all seek to win matches and improve their league standing. Fifth, the sample is large: we observe over 5,000 club-seasons between 1952 and 2022, providing statistical power to detect modest effects.

Sixth, we can track outcomes over a ten-year horizon, allowing us to examine both the short- and medium-run effects of windfalls.

Our main finding is that financial windfalls are associated with substantial improvements in league performance. Treated clubs climb approximately 18 league places over ten years—roughly three-quarters of a full division—and are 10.6 percentage points less likely to be relegated, with both effects highly significant at every horizon. Among clubs for which promotion is possible (those below the top division), treated clubs are 3 to 7 percentage points more likely to achieve promotion. We estimate linear probability models and linear regressions with a binary treatment indicator, reporting heteroskedasticity-consistent standard errors throughout.

We probe this finding in several ways. First, we decompose the effect by match outcome, distinguishing between clubs that win and clubs that lose their FA Cup match. This is important because winning an upset victory against a top-division club could generate psychological momentum, media attention, and supporter enthusiasm that are distinct from the pure financial windfall. We find that clubs that *lose* their match—receiving money without any confounding boost from victory—are significantly protected against relegation, with the effect growing from 2 percentage points at one year to nearly 9 percentage points at ten years. This provides a more direct test of the financial mechanism: losing clubs receive the windfall without any psychological, reputational, or competitive boost, yet their league position improves substantially over the following decade.

Second, we investigate whether organizational capacity moderates the treatment effect. If the returns to a financial windfall depend on an organization’s pre-existing infrastructure—its personnel, facilities, commercial operations, and strategic capacity—then better-resourced clubs should benefit more. We use stadium capacity as a proxy for this organizational capacity and find that the windfall’s effect on upward mobility through the league pyramid is strongly moderated by it: clubs with larger stadiums climb significantly more places, with the interaction significant at every horizon over the ten-year window. This moderating effect applies to upward movement but not to protection against relegation, suggesting that a financial buffer alone can prevent decline, while converting money into competitive improvement requires the organizational infrastructure to put it to use.

These results contribute to several literatures. For organizational theory and management, we provide credible causal evidence that financial resources improve performance—but that the effect on upward mobility is conditional on organizational capacity, while protection against decline requires only a financial buffer. This extends to financial resources a finding that Cohen and Levinthal (1990) established for knowledge: that an organization’s ability to benefit from new inflows depends on its prior complementary assets. It also bears on the debate over whether organizational slack is beneficial or wasteful (Nohria and Gulati 1996) and to the resource-based view’s foundational question of whether resources translate into competitive advantage (Barney 1991). For the corporate finance literature on windfalls, our results suggest that not all unexpected cash inflows lead to waste or agency problems—a finding that complements Blanchard et al. (1994) by studying a much larger sample with random treatment assignment. For the study of money and performance in competitive settings, our setting provides stronger identification than is typically available in the campaign

finance (Levitt 1994) or sports economics (Szymanski 2000) literatures.

2 Related Literature

Our paper sits at the intersection of several literatures that study the relationship between financial resources and organizational outcomes. As we see it, there are four main strands: the management literature on financial resources and performance; the evidence on cash windfalls from corporate and public finance; the question of organizational capacity and the conditions under which resources help; and the study of money and performance in competitive settings.

2.1 Financial Resources and Organizational Performance.

A central question in management and organizational theory is whether financial resources contribute to organizational success. The resource-based view of the firm (Barney 1991) argues that sustained competitive advantage derives from resources that are valuable, rare, imperfectly imitable, and non-substitutable. While financial capital is typically considered fungible and therefore not a source of *sustained* advantage, the question of whether exogenous increases in financial resources improve performance—even temporarily—is surprisingly difficult to answer.

The organizational slack literature addresses a closely related question: what happens when organizations possess resources beyond what is strictly needed for current operations? Cyert and March (1963) introduced the concept of organizational slack in their behavioral theory of the firm, arguing that slack serves as a buffer that facilitates organizational adaptation and conflict resolution. Subsequent work has debated whether slack is beneficial or harmful. Nohria and Gulati (1996) propose an inverted-U relationship: too little slack constrains experimentation and adaptation, while too much slack breeds complacency and wasteful expenditure. This tension is central to our paper. A financial windfall represents a sudden increase in organizational slack, and whether it improves or harms performance may depend on the organization’s capacity to deploy it productively.

The fundamental challenge in this literature is endogeneity. Organizations that perform well generate more resources, and unobserved managerial quality simultaneously drives both resource acquisition and performance. As a result, the extensive correlational evidence linking financial inputs to organizational outcomes is difficult to interpret causally. Our paper addresses this challenge directly by exploiting a setting in which financial resources are randomly assigned.

2.2 Cash Windfalls in Corporate and Public Finance.

The most direct existing evidence on the causal effect of money on organizational behavior comes from studies of cash windfalls. In corporate finance, Blanchard et al. (1994) study eleven firms that received large cash windfalls from winning or settling lawsuits. They find that managers used these windfalls to pursue empire-building and entrenchment rather than maximizing shareholder value—evidence consistent with the agency costs of free cash flow theorized by Jensen (1986). This landmark

study suggests that unexpected money can worsen organizational outcomes by exacerbating agency problems between managers and shareholders.

In public finance, parallel findings emerge from studies of government windfalls. Brollo et al. (2013) show that Brazilian municipalities receiving larger-than-expected federal transfers experience increased corruption and lower-quality political candidates—a “political resource curse” driven by weakened electoral accountability. Caselli and Michaels (2013) find that oil royalties in Brazilian municipalities lead to higher reported spending but little improvement in actual public goods provision, with the discrepancy suggesting waste or misappropriation. Berset and Schelker (2020) document that windfall revenues lead Swiss municipalities to increase spending and accumulate debt. At the individual level, Carlson et al. (2015) find that NFL players—who earn a median of \$3.2 million over careers averaging six years—file for bankruptcy at rates comparable to the general population, with the size of career earnings having no detectable effect on post-retirement financial distress. The failure to smooth consumption persists even for the highest earners.

The consistent finding across all of these settings—firms, governments, individuals—is negative: unexpected or short-lived money tends to be wasted, misused, or inadequately managed. However, these settings differ from ours in ways that may explain the negative results. The corporate windfall studies involve very small samples and imperfect identification. The public finance studies involve political actors with weak accountability and diffuse objectives. The individual-level evidence involves people without organizational structures to channel resources. Our setting differs in that treatment is randomly assigned across a large sample of organizations that operate in a competitive environment with clear, measurable objectives.

2.3 Organizational Capacity and the Conditions Under Which Resources Help.

If financial resources do not automatically improve performance, what determines whether they help or harm? A structural insight from several literatures is that the returns to a new external input depend on the recipient’s prior stock of complementary assets.

Cohen and Levinthal (1990) established this principle in the domain of knowledge and innovation, introducing the concept of absorptive capacity—the ability of an organization to recognize the value of new external information, assimilate it, and apply it to commercial ends. Their key argument is that this ability depends on the organization’s prior related knowledge and develops cumulatively: firms that have invested in R&D are better positioned to benefit from external knowledge spillovers. Zahra and George (2002) extended the framework, reconceptualizing absorptive capacity as a set of organizational routines and processes—acquisition, assimilation, transformation, and exploitation—that constitute a dynamic capability.

Although Cohen and Levinthal’s framework was developed for knowledge absorption specifically, the underlying logic—that pre-existing organizational infrastructure determines the returns to new inflows—applies more broadly. In the dynamic capabilities literature, Teece et al. (1997) argue that an organization’s ability to “integrate, build, and reconfigure internal and external competences” determines whether it can convert new resources into competitive advantage. An organization that

receives an unexpected cash windfall can only benefit from it if it has the personnel, infrastructure, and strategic capacity to identify and execute worthwhile investments. Without such organizational capacity, the windfall may be wasted on poorly conceived expenditures, consumed by agency costs, or simply dissipated.

This insight parallels findings in the resource curse literature, where institutional quality has been shown to moderate the effect of resource windfalls on economic outcomes. Mehlum et al. (2006) demonstrate that the resource curse is concentrated in countries with “grabber-friendly” institutions, while countries with “producer-friendly” institutions manage resource wealth effectively. Robinson et al. (2006) develop a theoretical model in which the effect of resource booms depends on the quality of political institutions. These findings at the country level suggest a general principle: the effect of financial windfalls is conditional on the recipient’s organizational capacity to deploy them.

Our paper tests this directly in a setting with random treatment assignment. By interacting the windfall treatment with a proxy for organizational capacity (home stadium capacity), we can estimate whether the effect of money on performance is moderated by the organization’s pre-existing infrastructure—bringing the absorptive capacity logic into a domain where the inflow is cash rather than knowledge.

2.4 Money and Performance in Competitive Settings.

The question of whether money causes success has been extensively studied in two competitive domains: electoral politics and sports. In the campaign finance literature, Jacobson (1978) documents a strong positive correlation between campaign spending and vote share, but the causal interpretation is contested. Levitt (1994) exploits repeat challengers in U.S. House races to control for candidate quality and finds much smaller spending effects, while Gerber (1998) uses instrumental variables in Senate races and finds larger effects. The debate remains unresolved, largely because of the difficulty of finding exogenous variation in campaign resources.

In sports economics, the relationship between financial expenditure and team performance is well established but similarly complicated by endogeneity. Szymanski and Smith (1997) examine the English football industry and find a strong relationship between wage expenditure and league position, but the direction of causality is ambiguous: successful clubs generate more revenue, enabling higher wages, which in turn may produce further success. Szymanski (2000) uses wage data to test for discrimination in English football, further demonstrating the tight link between payroll and performance while highlighting the identification challenges inherent in the relationship. More recently, Kleven et al. (2013) use the European football labor market to study how tax rates drive international player migration, demonstrating both the credibility of football data for studying financial incentives and the rigid labor demand conditions—fixed squad sizes, displacement of domestic players by foreigners—under which clubs compete for talent.

Our FA Cup design provides cleaner identification than is typically available in either of these literatures. The random draw ensures that the financial treatment is orthogonal to club quality, and the large sample and long time horizon provide the statistical power to detect effects and explore

heterogeneity.

3 Institutional Background: The FA Cup

The Football Association Challenge Cup, commonly known as the FA Cup, is the oldest national football competition in the world, first played in the 1871–72 season. It is an annual knockout tournament open to all eligible clubs in the English football league system, from the Premier League (the top division) down to level nine and below. The competition features multiple qualifying rounds before the “proper” rounds begin. In recent decades, over 700 clubs have entered annually. The Cup runs concurrently with the league season—in which the same clubs compete for divisional championships through a separate year-long round-robin—but the two competitions are entirely independent: cup results have no effect on league standings, and the league outcomes we study (position, promotion, relegation) are determined solely by league play.

To understand how the FA Cup generates financial windfalls—and why promotion is a meaningful outcome measure—it is necessary to describe the structure of English football and the role of the cup competition within it.

3.1 The English Football Pyramid.

English football is organized as a hierarchical pyramid of interconnected leagues, with promotion and relegation linking each level. Over 7,000 clubs compete across the system, making it one of the largest and most stratified league structures in world sport.

The top four levels constitute professional football. The Premier League (Level 1) contains 20 clubs and is among the wealthiest sports leagues in the world. Below it sit the three divisions of the English Football League (EFL): the Championship (Level 2, 24 clubs), League One (Level 3, 24 clubs), and League Two (Level 4, 24 clubs). Below the EFL, the National League (Level 5) and its regional feeders extend down through progressively more localized divisions. The names and branding of these tiers have changed over time—most notably, the top tier was known as Division 1 until the creation of the Premier League in 1992—but the hierarchical structure has remained stable throughout our sample period. We refer to divisions by number (Division 1 through Division 4) regardless of their official names in a given era.

Within each division, every club plays each rival twice over the course of a season—once at home and once away—earning three points for a win and one for a draw (two points for a win before the 1981–82 season). The final league table, ranked by total points, determines the divisional champion; winning the top division is the most coveted achievement in English football, comparable in stature to winning the World Series or Super Bowl. What makes the English system distinctive, however, is promotion and relegation: at the end of each season, the top clubs in each division move up to the tier above while the bottom clubs drop down. In the Premier League, the bottom three clubs are relegated to the Championship. In the Championship, the top two clubs earn automatic promotion to the Premier League, with a third promotion place determined by a playoff among the

clubs finishing third through sixth. Similar mechanisms operate at each level of the pyramid. This means that a club’s divisional placement is a direct, objective measure of competitive success, and promotion represents a discrete, unambiguous improvement in organizational standing. Figure 1 illustrates this structure, showing the number of clubs and the promotion and relegation flows at each level.

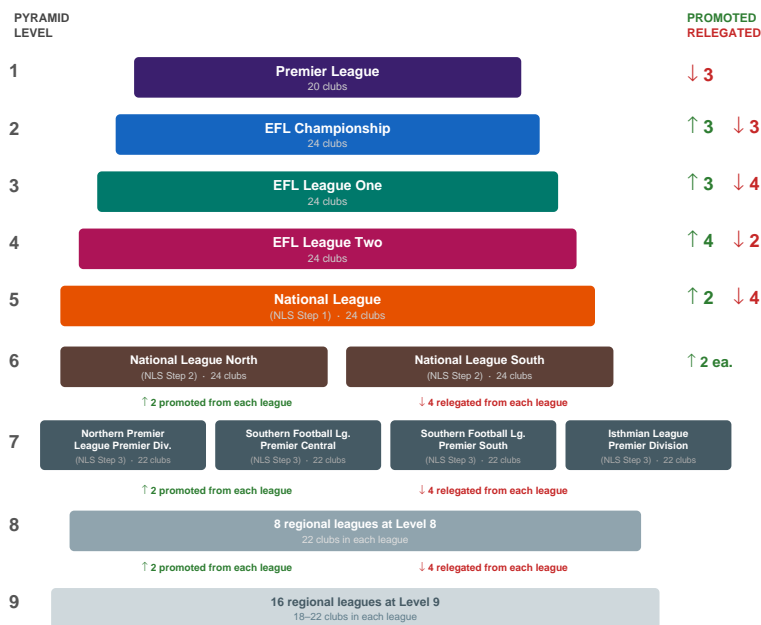


Figure 1: The English football pyramid. Each level shows the number of clubs and the number promoted (up arrows) and relegated (down arrows) each season. The top four levels constitute professional football; below Level 5 the system branches into parallel regional leagues.

The financial consequences of promotion are substantial. Moving up a single division brings increased gate receipts, higher broadcasting revenue, and greater commercial income. Promotion from the Championship to the Premier League, for example, is worth tens of millions of pounds in additional annual revenue. Even at lower levels, promotion brings meaningfully higher revenues. Because promotion is both objectively measured and economically consequential, it serves as an ideal outcome variable for studying whether financial windfalls improve organizational performance.

3.2 The FA Cup Competition.

The competition is structured as a single-elimination tournament: clubs are paired in each round, and the winner advances. The critical feature for our research design is the **third round proper**, the stage at which the 20 Premier League and 24 Championship (second-division) clubs enter the competition, joining the survivors from lower divisions. At this stage, approximately 64 clubs are drawn randomly into pairings, with one club designated as the home team.

3.3 The Random Draw.

The third-round draw is conducted publicly and is genuinely random: numbered balls representing each club are placed into a pot and drawn sequentially, with the first ball drawn in each pair designated as the home team. There is no seeding, no geographical separation, and no protection for any club. This means that a non-league club can be drawn away at a Premier League club's stadium, or two lower-league clubs can be drawn against each other.

3.4 Gate Receipts and the Financial Windfall.

Under FA rules, gate receipts from each match are shared between the two competing clubs after the FA takes its share and stadium costs are deducted. Specifically, for matches from the third round proper through the sixth round, 10% of net gate receipts (total ticket revenue less VAT and match-day expenses) is retained by the FA for a central prize pool, and the remaining 90% is split equally between the two clubs—45% each. Where one of the clubs is from outside the Football League, that club receives a larger share: 50% versus 40% for the league club. In either case, the away club receives its share of gate receipts from a match played in the home club's stadium.

This sharing rule is critical to our research design. When a small club is drawn to play away at a large club's stadium, the match is played in a much larger venue than the small club's home ground, and the small club receives its share of the resulting gate receipts. For clubs in the lower leagues and below, it is often more financially attractive to be drawn *away* against a top club than to host the match at home, precisely because the away draw delivers a share of revenues from a far larger stadium.

The magnitude of these windfalls can be transformative. According to a Deloitte analysis of FA Cup finances, the competition delivered £643 million in prize money, television payments, and gate receipts to participating clubs over the decade from 2001–02 to 2010–11, with gate receipts alone totaling £413 million. For clubs at the top of the pyramid, FA Cup revenues are a small fraction of total income—Stoke City's run to the sixth round in 2009–10 generated approximately £2.1 million, roughly 4% of the club's annual revenue. But for clubs further down the pyramid, the same competition can be a financial lifeline. When Havant & Waterlooville, a Conference South club, reached the fourth round in 2007–08 and were drawn away at Liverpool's Anfield, the run generated approximately £600,000—some 70% of the club's estimated annual revenue. Leyton Orient's run to the fifth round in 2010–11, which included a tie at Arsenal's Emirates Stadium, generated £1.4 million, approximately 30% of the League One club's annual income.

Consider a further example that illustrates the link between windfalls and subsequent performance. In the 2005–06 third round, Burton Albion (then a non-league club with average home attendance of approximately 1,724) were drawn at home against Manchester United, forced a replay, and then traveled to Old Trafford (then approximately 68,000 capacity, mid-way through its 2005–06 expansion to 76,212). The two matches generated reportedly close to £1 million in combined revenue for Burton—a transformative sum for a club at that level. Burton's manager Nigel

Clough later reflected that the money “paid for the ground... gave them the foundation to go on and do what they did.” Burton were promoted to the Football League three seasons later.

In addition to gate receipts, clubs earn prize money for each round they win (approximately £67,500 per third-round victory in 2011–12) and television fees if their matches are selected for live broadcast (approximately £123,500 per televised club in the third round). These additional revenue streams further increase the financial windfall from a deep cup run, though gate receipts from matches at large stadiums remain the primary source of the treatment we exploit.

4 Research Design

4.1 Treatment Definition.

Our treatment exploits the random assignment of FA Cup third-round matchups. We define treatment using a binary indicator:

$$\text{BinTreat}_i = \begin{cases} 1 & \text{if club } i \text{ is drawn away at a club from a higher division} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

A treated club is one that is randomly assigned to play as the away team at the stadium of a club from a higher division. Such a club receives a share of gate receipts from a match played in a larger-than-usual venue, generating a positive financial shock. Clubs that are drawn at home, that face a club from the same or a lower division, or that are the higher-division club in a cross-division pairing serve as the control group. Importantly, this comparison also isolates the financial channel from non-financial effects of the match. A small club drawn at home against a top-division opponent receives the same media exposure, prestige, and scouting attention as one drawn away—the only difference is who keeps the gate receipts from the larger stadium. Because home draws are in the control group, the estimate captures the financial windfall net of any non-financial effects of facing a prominent opponent.

We also consider a continuous treatment measure that captures the magnitude of the financial windfall, defined as the difference in divisional rank between the opponent and the focal club.

4.2 Identification Strategy.

The key identifying assumption is that the FA Cup draw is random, which it is by construction. Conditional on entering the third round proper, the identity of the opponent and the home/away designation are determined by a random draw. This means that treatment assignment is independent of any club characteristics—current form, financial situation, managerial quality, or any other potential confound.

Formally, let $Y_{i,t+k}$ denote the outcome for club i at horizon k years after the FA Cup draw in

season t . Our estimating equation is:

$$Y_{i,t+k} = \alpha_k + \beta_k \cdot \text{BinTreat}_{it} + \varepsilon_{it} \quad (2)$$

where β_k is the causal effect of the financial windfall on the outcome at horizon k . Because treatment is randomly assigned, β_k is identified without the need for controls, fixed effects, or instrumental variables. We report heteroskedasticity-consistent standard errors (HC1; White 1980) throughout, treating each FA Cup draw as an independent randomization. Because the same club can appear in multiple draws over our 70-year sample, one might alternatively cluster standard errors by club to account for within-club serial correlation in outcomes. We report all main results under club-clustered standard errors in Appendix B and find that conclusions are unchanged.

Why we do not include division fixed effects. A natural question is whether one should condition on the club’s starting division. We deliberately do not, for three reasons. First, randomization makes it unnecessary for identification: the draw is orthogonal to division and all other pre-treatment characteristics, so there is no confounding to remove. Second, the treatment probability is structurally related to division—lower-division clubs face more potential higher-division opponents and are therefore more likely to be treated—so division fixed effects alter the estimand rather than purify it. The pooled estimate $\hat{\beta}_k$ recovers the average effect of a randomly assigned windfall across the population of clubs that enter the FA Cup, which is the policy-relevant parameter. Division fixed effects would instead recover a variance-weighted average of within-division effects, downweighting precisely the divisions where windfalls are proportionally largest and most consequential. Third, the windfall’s financial significance varies across divisions—the same gate receipt represents a transformative shock for a Division 4 club but a rounding error for a Championship side—so conditioning on division absorbs economically meaningful variation in treatment intensity. The attenuation that results from adding division fixed effects is therefore expected and informative: it reflects heterogeneous treatment effects across divisions, not omitted variable bias. We return to this point empirically in Section 6.

Validation: pre-treatment trajectories. Because treatment probability varies by division, treated clubs are disproportionately drawn from lower divisions. Lower-division clubs have mechanically different pre-treatment trajectory distributions—more room to have been recently promoted or relegated—so treated and control groups differ on pre-treatment trajectory measures in unconditional comparisons (Appendix Table 6). This does not reflect a failure of randomization; it reflects the compositional variation in who receives the largest windfalls. Conditional on division, all nine trajectory variables we examine—including tier changes, prior promotion and relegation indicators, and a five-year trajectory slope—are precisely balanced (all $p > 0.19$). Moreover, pre-treatment trajectory does not moderate the treatment effect: the interaction between treatment and trajectory is insignificant at every horizon for all trajectory measures (Appendix Table 7). This confirms that the compositional difference does not bias our estimates.

4.3 Outcome Measures.

We measure the effect of financial windfalls on three outcomes that capture different dimensions of movement through the league pyramid, each evaluated at horizons $k = 1, 2, \dots, 10$ years following the FA Cup draw.

Our first outcome is *places climbed*: the change in a club’s global league position between baseline and horizon k , where global position is defined as the sum of teams in all higher divisions plus rank within the club’s own division. This continuous measure captures upward and downward movement within and across divisions.

Our second outcome is *relegation*: whether a club occupies a lower division at horizon k than at baseline. This binary indicator captures the risk of decline.

Our third outcome is *promotion*: whether a club has moved to a higher division by horizon k (cumulative). Because promotion is mechanically impossible for top-division clubs, we restrict the promotion analysis to clubs in Divisions 2 and below.

We estimate Equation (2) separately for each horizon k , producing a trajectory of treatment effects over the ten-year window.

5 Data

We compile a comprehensive dataset of all FA Cup third-round proper matches from 1952—the earliest season for which we can construct complete league-position histories—to 2022, covering 70 seasons of the competition. For each match, we record the identities of the two clubs, the home/away designation, the league and divisional affiliation of each club, the match result, and the stadium and capacity information. Stadium capacities reflect historical values for each era, accounting for stand construction, safety-certificate requirements under the Safety of Sports Grounds Act (1975), further restrictions following the Popplewell inquiry (Popplewell 1986), and the post-1990 conversion to all-seater stadiums mandated by the Taylor Report (Taylor 1990), rather than a single fixed value per ground (see Appendix E for details). We then track each club’s league position (division and rank) for up to ten years following each FA Cup appearance.

The resulting dataset contains 5,349 club-season observations with non-missing treatment status. Table 1 presents summary statistics.

Table 1: Summary statistics

Variable	Mean	Std. Dev.	Min	Max
<i>Treatment and club characteristics</i>				
Binary treatment (BinTreat)	0.175	0.380	0	1
Continuous treatment (ConTreat)	0.367	1.033	0	9
Home indicator	0.542	0.498	0	1
Won match	0.360	0.480	0	1
Division	2.47	1.70	1	10
Home stadium capacity	30,240	15,444	1,000	76,000
<i>Outcomes (full sample, $N = 5,349$)</i>				
Places climbed within 1 year	0.3	12.4	-92	97
Places climbed within 5 years	-1.7	24.1	-103	101
Places climbed within 10 years	-2.3	28.8	-105	100
Relegated within 1 year	0.103	0.303	0	1
Relegated within 5 years	0.251	0.434	0	1
Relegated within 10 years	0.266	0.442	0	1
<i>Outcomes (Divisions 2+, $N = 3,509$)</i>				
Promoted within 1 year	0.092	0.290	0	1
Promoted within 5 years	0.189	0.392	0	1
Promoted within 10 years	0.183	0.387	0	1

Notes: Summary statistics for all club-season observations in the FA Cup third round proper, 1952–2022. Binary treatment equals one if a club is drawn away at a higher-division club. Continuous treatment measures the divisional difference between opponent and focal club. Places climbed and relegation are reported for the full sample ($N = 5,349$); promotion is reported for Divisions 2 and below ($N = 3,509$) because promotion is mechanically impossible for top-division clubs. Sample sizes for outcome variables vary by horizon due to missing standings data at longer horizons.

6 Results

6.1 Main Results: Movement Through the Pyramid.

We begin with outcomes that are well-defined for all clubs in the sample, including top-division clubs: league position (places climbed) and relegation. We then turn to promotion, which requires restricting the sample to clubs below the top division.

Places climbed. Figure 2 presents the effect of the windfall on global league position. Each point represents the estimated coefficient $\hat{\beta}_k$ from Equation (2), where the dependent variable is the change in a club’s global league position (the sum of teams in all higher divisions plus rank within its own division) between the baseline and horizon k . Positive values indicate upward movement.




Figure 2: Effect of financial windfall on global league position (places climbed). Positive values indicate upward movement. Each point shows the coefficient from a linear regression of places climbed at horizon k on the binary treatment indicator. Error bars show \pm one standard error.

The treatment effect is positive and highly significant ($p < 0.001$) at every horizon. Treated clubs climb approximately 3.3 league places within one year, and the effect grows steadily to 17.8 places at ten years—roughly three-quarters of a full division. The monotonic growth of the effect is consistent with windfalls setting off a sustained trajectory of investment and improvement, not a one-time bump. Because places climbed is a continuous measure that captures within-division improvement, cross-division transitions, and their interaction, this result cannot be attributed to compositional effects: a club that climbs 17 league places has genuinely moved up the table, regardless of where it started.

Relegation. Figure 3 presents the effect of the windfall on relegation—an indicator equal to one if a club occupies a lower division at horizon k than at baseline. The windfall significantly reduces the probability of relegation at every horizon, with effects of -3.4 percentage points at one year growing to -10.6 percentage points at ten years (all $p < 0.01$). The windfall does not only improve upward mobility—it also provides significant protection against relegation. This finding is consistent with the organizational slack literature’s characterization of excess resources as a buffer against decline (Cyert and March 1963).




Figure 3: Effect of financial windfall on relegation probability. Each point shows the coefficient from a linear probability model of relegation at horizon k on the binary treatment indicator. Negative values indicate protection against relegation. Error bars show \pm one standard error.

Table 2 presents the underlying regression estimates for both outcomes.

Because many clubs appear in the FA Cup draw more than once, our main estimates pool first and subsequent treatments. When we restrict the sample to each club’s first treatment—removing any contamination from prior windfalls—the effects roughly double in magnitude, with treated clubs climbing 38 places over ten years and relegation risk falling by 21 percentage points (Appendix Table 17). This is consistent with repeated treatments adding attenuation rather than inflation: a club’s first windfall is a clean shock, while subsequent windfalls arrive against a baseline already shaped by prior treatment.

Table 2: Main results: Effect of financial windfall on league position

Horizon	Panel A: Places Climbed				Panel B: Relegation			
	Coef.	SE	p	N	Coef.	SE	p	N
1 year	3.321***	(0.736)	<0.001	5,013	-0.034***	(0.010)	0.001	5,293
2 years	5.015***	(1.010)	<0.001	4,891	-0.045***	(0.012)	<0.001	5,266
3 years	7.906***	(1.136)	<0.001	4,790	-0.067***	(0.013)	<0.001	5,241
4 years	10.196***	(1.302)	<0.001	4,684	-0.074***	(0.014)	<0.001	5,216
5 years	11.620***	(1.445)	<0.001	4,589	-0.066***	(0.015)	<0.001	5,206
6 years	11.776***	(1.497)	<0.001	4,508	-0.080***	(0.015)	<0.001	5,200
7 years	13.977***	(1.676)	<0.001	4,420	-0.088***	(0.015)	<0.001	5,191
8 years	14.825***	(1.645)	<0.001	4,325	-0.096***	(0.014)	<0.001	5,181
9 years	16.704***	(1.694)	<0.001	4,200	-0.102***	(0.014)	<0.001	5,179
10 years	17.834***	(1.752)	<0.001	4,111	-0.106***	(0.014)	<0.001	5,171

Notes: Each row reports the coefficient on the binary treatment indicator from a separate regression. Panel A: dependent variable is the change in global league position (places climbed) at horizon k ; positive values indicate upward movement. Panel B: dependent variable is an indicator for relegation (occupying a lower division) at horizon k . Sample sizes vary by horizon because places climbed requires non-missing rank data at horizon k . Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

6.2 Promotion.

Promotion is a binary threshold: a club either crosses it or does not. It is also the most direct measure of the outcome that matters most to clubs. However, because the top division has no tier above it, promotion is mechanically impossible for top-division clubs: these observations contribute only zeros to the promotion outcome and inflate the estimated treatment effect. We therefore estimate promotion regressions on the sample of clubs in Divisions 2 and below ($N = 3,509$), for whom promotion is a non-degenerate outcome.

Figure 4 and Table 3 present the results. The treatment effect on cumulative promotion is positive and statistically significant at every horizon, growing from 3.3 percentage points at one year ($p = 0.019$) to 6.5 percentage points at ten years ($p < 0.001$). These effects are smaller than those from the full sample—which is expected, since excluding top-division clubs removes the observations where the windfall is proportionally largest relative to a club’s revenues—but they remain economically meaningful. A 6.5 percentage-point increase in the ten-year cumulative probability of promotion represents a meaningful improvement for a lower-division club.

Division composition. As discussed in Section 4, we do not include division fixed effects in our main specifications. For the places climbed and relegation results, division fixed effects are unnecessary for identification (treatment is randomly assigned) and would absorb cross-division variation in treatment intensity that is part of the causal effect. For promotion, the relevant compositional issue is not fixed effects but the mechanical zero problem: top-division clubs cannot be promoted, so including them inflates the treatment coefficient. Restricting to Divisions 2 and below resolves this directly. When we additionally add division fixed effects within this restricted sample, the promo-

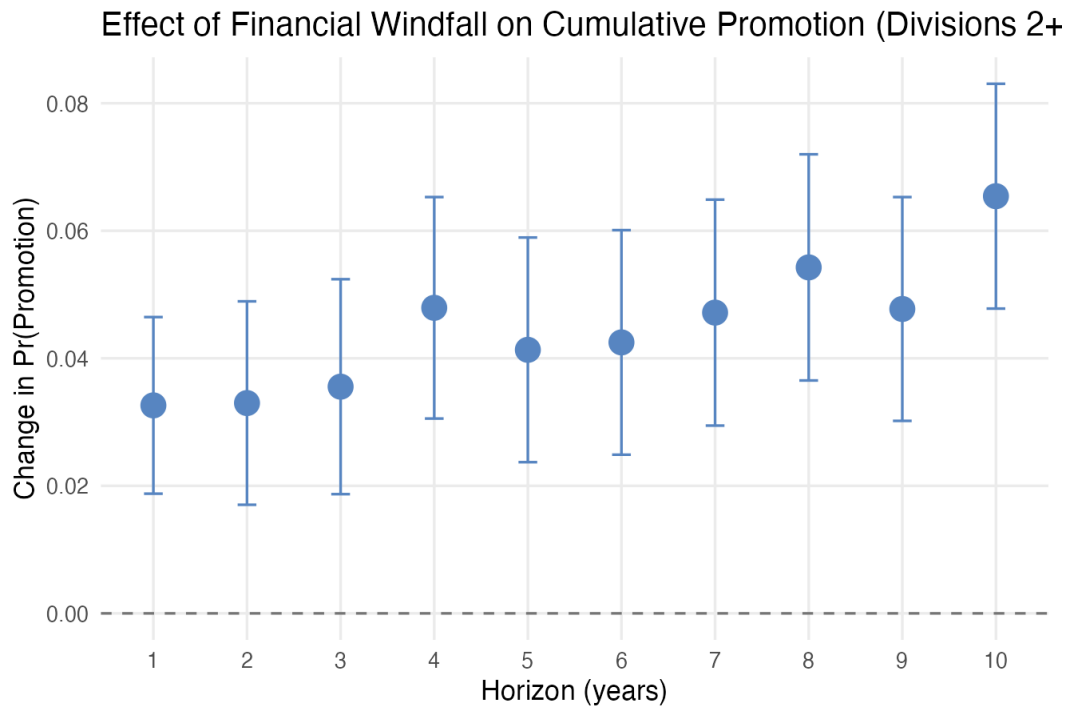


Figure 4: Effect of financial windfall on cumulative promotion probability (Divisions 2 and below). Each point shows the coefficient from a linear probability model of cumulative promotion at horizon k on the binary treatment indicator. Error bars show \pm one standard error. Sample restricted to clubs for whom promotion is a non-degenerate outcome ($N = 3,509$).

Table 3: Effect of financial windfall on promotion (Divisions 2 and below)

Horizon	Coef.	SE	p
1 year	0.033*	(0.014)	0.019
2 years	0.033*	(0.016)	0.039
3 years	0.036*	(0.017)	0.035
4 years	0.048**	(0.017)	0.006
5 years	0.041*	(0.018)	0.019
6 years	0.042*	(0.018)	0.016
7 years	0.047**	(0.018)	0.008
8 years	0.054**	(0.018)	0.002
9 years	0.048**	(0.018)	0.007
10 years	0.065***	(0.018)	<0.001
N		3,509	

Notes: Each row reports the coefficient on the binary treatment indicator from a separate linear probability model. Dependent variable is cumulative promotion (promoted to a higher division by horizon k). Sample restricted to clubs in Divisions 2 and below, for whom promotion is a non-degenerate outcome. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

tion coefficient attenuates further, consistent with heterogeneous treatment effects across divisions where the windfall is proportionally larger for clubs in lower divisions.

6.3 Robustness: Rolling Estimates Over Time.

To assess whether the treatment effect is stable over the seven decades in our sample, we estimate rolling five-year average treatment effects. Figure 5 plots the five-year rolling average of the yearly treatment coefficient over the sample period. The trajectory of the treatment effect over time is not merely a robustness exercise; it provides substantive insight into the mechanism through which windfalls affect organizational performance.

The treatment effect is small and near zero in the early years of our sample, during the 1950s. It rises through the 1960s and 1970s, peaks sharply in the mid-to-late 1980s, and then gradually declines from the mid-1990s onward, settling at a lower but still positive level by the 2010s. This temporal pattern closely tracks the evolving role of gate receipts in English football club finances, providing additional evidence that the mechanism driving our results is financial rather than psychological.

The low early estimates (1950s). In the immediate postwar period, English football experienced record attendance levels, with aggregate Football League attendances peaking at 41.2 million in 1948–49 (Dobson and Goddard 2001). Although attendances were beginning their long decline by the 1950s, clubs at all levels of the pyramid still drew substantial crowds relative to later decades. In this environment, the marginal financial impact of an FA Cup windfall—while positive—was relatively modest compared to a club’s normal gate revenue. The treatment, in other words, was a smaller proportional shock.




Figure 5: Five-year rolling average of the treatment effect over time. Each point is the rolling mean of the yearly treatment coefficient, 1952–2022.

The rising effect (1960s–1980s). From the mid-1960s onward, English football entered a prolonged period of declining attendance and financial distress. Aggregate Football League attendances fell by approximately 60 percent between the postwar peak and the mid-1980s nadir of 16.5 million in 1985–86 (Szymanski 2003, Dobson and Goddard 2001). The causes were multiple: the rise of football hooliganism, deteriorating stadium infrastructure, competition from television and other leisure activities, and broader economic difficulties including the recessions of the mid-1970s and early 1980s. For lower-league clubs, which depended almost entirely on gate receipts, this period was one of severe financial constraint. In this context, an FA Cup windfall from playing at a large top-division stadium represented an increasingly large proportional shock to a club’s finances. A single cup tie at a 40,000-seat ground could deliver revenue equivalent to months of normal operations for a club whose home attendance had dwindled to a few thousand.

The peak (mid-to-late 1980s). The treatment effect reaches its maximum around 1985–1990, a period that coincides with English football’s deepest crisis. The Heysel Stadium disaster in 1985 led to a five-year ban on English clubs in European competition. The Bradford City stadium fire in 1985 and the Hillsborough disaster in 1989 exposed the dangerous state of the country’s ageing grounds and prompted expensive safety-mandated stadium improvements following the Taylor Report (Taylor 1990). Lower-league clubs were particularly cash-strapped during this period, and the financial impact of a cup windfall would have been at its maximum relative to clubs’ baseline revenues. Moreover, with fewer alternative revenue streams available, the windfall represented one

of the only mechanisms through which a small club could obtain a substantial, discrete financial injection.

The decline in the Premier League era (1992 onward). The creation of the FA Premier League in 1992 and the accompanying television deal with BSkyB—worth £304 million over five years—fundamentally restructured the economics of English football (Szymanski and Kuypers 1999). Television revenue rapidly came to dominate club finances at the top of the pyramid, and the financial gap between the Premier League and the lower divisions widened dramatically. By the 2020s, the average Premier League club received approximately £100 million per season in broadcasting revenue alone, while the average League One or League Two club received solidarity payments of only a few hundred thousand pounds. This growing gap has two implications for our treatment effect. First, the windfall from an FA Cup tie, while still significant in absolute terms, represents a smaller proportional shock for lower-league clubs that now have access to more diverse (if still modest) revenue streams, including their own broadcasting deals, commercial sponsorships, and in some cases wealthy individual owners. Second, the cost of improving league position has escalated substantially. Wage inflation across all divisions means that the competitive advantage purchasable with a given windfall has shrunk; what might have funded two or three quality signings in the 1980s may cover only a fraction of one player’s annual salary in the 2020s.

Interpretation. The temporal variation in the treatment effect speaks directly to our central question. If the windfall operated primarily through non-financial channels—psychological momentum, media exposure, or supporter enthusiasm from an FA Cup appearance—there is no obvious reason the effect should track the relative financial importance of gate receipts over time. A “cup magic” story would predict, if anything, a stable or growing effect as media coverage of FA Cup upsets intensified in the television era. Instead, the effect rises and falls with the financial salience of gate receipt windfalls relative to clubs’ overall revenues, consistent with a financial mechanism. This temporal pattern also reinforces our organizational capacity finding (Section 7.2). The post-1992 era created a football economy in which the financial gulf between divisions is so large that even a substantial gate receipt windfall may be insufficient to bridge it without considerable pre-existing organizational capacity. The declining average treatment effect in recent decades may partly reflect the growing difficulty of converting a fixed-size financial shock into league improvement when the financial requirements of competitive success have escalated dramatically.

7 Mechanisms and Moderators

7.1 Winning versus Losing: Disentangling Money from Momentum.

A potential concern with our identification strategy is that being drawn against a top-division club might affect a smaller club’s subsequent performance through channels other than money. An upset victory in the FA Cup could generate psychological momentum, media attention, increased

supporter enthusiasm, and a boost to the club’s reputation—all of which could independently affect league performance. This concern echoes a persistent challenge in the campaign finance literature, where correlations between spending and electoral success may reflect candidate quality rather than the causal effect of money (Jacobson 1978, Levitt 1994).

To address this, we decompose the treatment effect by match outcome, estimating separate effects for clubs that win and clubs that lose their FA Cup match. If the treatment effect operates primarily through the financial windfall rather than through the psychological boost of victory, we would expect to find effects among the losers—who receive the money without the confounding psychological effect—that are comparable to those among winners. We use relegation as the dependent variable for this decomposition because it is well-defined for all clubs in the sample, avoiding the need to restrict to a subsample.

Upset victories are uncommon: only 230 of the 935 treated clubs (24.6%) won their match, while 705 lost (including drawn matches resolved by a replay—a rematch typically held at the other club’s ground). This is unsurprising given that treated clubs are, by definition, lower-division sides playing away at higher-division opponents. The losers subsample is therefore substantially larger and yields more precisely estimated coefficients.

Figure 6 presents the results. The key finding is that *losing* clubs—those that received the financial windfall without any confounding psychological boost from victory—show large and statistically significant protection against relegation. The loser treatment effect grows from -2.6 percentage points at a one-year horizon to -9.1 percentage points at ten years, and is significant at every horizon (Table 4). This provides strong evidence that the financial channel is operative: clubs that received money but experienced no upset victory, no media celebration, and no psychological momentum still show clear improvements in subsequent league position.

The loser effect grows steadily over the decade. Unlike a psychological boost, which would be expected to dissipate quickly, the financial effect compounds over time as clubs reinvest windfall resources into players, staff, and infrastructure. The loser effect is comparable in magnitude to the winner effect at longer horizons—indeed, at 10 years, losers (-9.1 pp) show a larger point estimate than winners (-6.0 pp)—consistent with the interpretation that the durable effect is driven by money rather than momentum.

Winners show a larger effect at year 1 (-6.6 pp vs. -2.6 pp), but this gap is difficult to attribute cleanly to psychology. Clubs that win their FA Cup match also advance to the next round, earning additional prize money and potentially another lucrative gate receipt from a subsequent tie—so even the year 1 winner premium may partly reflect additional financial treatment rather than pure momentum.

The pattern is not limited to relegation. When we use places climbed as the dependent variable, losers climb 14 places over ten years—approximately two-thirds of the winner effect of 22 places—and the loser effect is highly significant at every horizon (Appendix Table 8). The convergence between winners and losers at longer horizons is again consistent with a financial mechanism that compounds over time, independent of any psychological boost from victory.

This decomposition isolates the financial channel more cleanly than is typically possible in windfall studies. In corporate windfall studies (Blanchard et al. 1994), the lawsuit itself may signal firm characteristics; in political windfall studies (Brollo et al. 2013), transfers are correlated with municipal attributes. Here, the losers subsample isolates the pure financial channel: these clubs received money but no psychological, reputational, or competitive boost from victory, yet they are significantly less likely to be relegated and climb substantially through the league pyramid over the following decade.

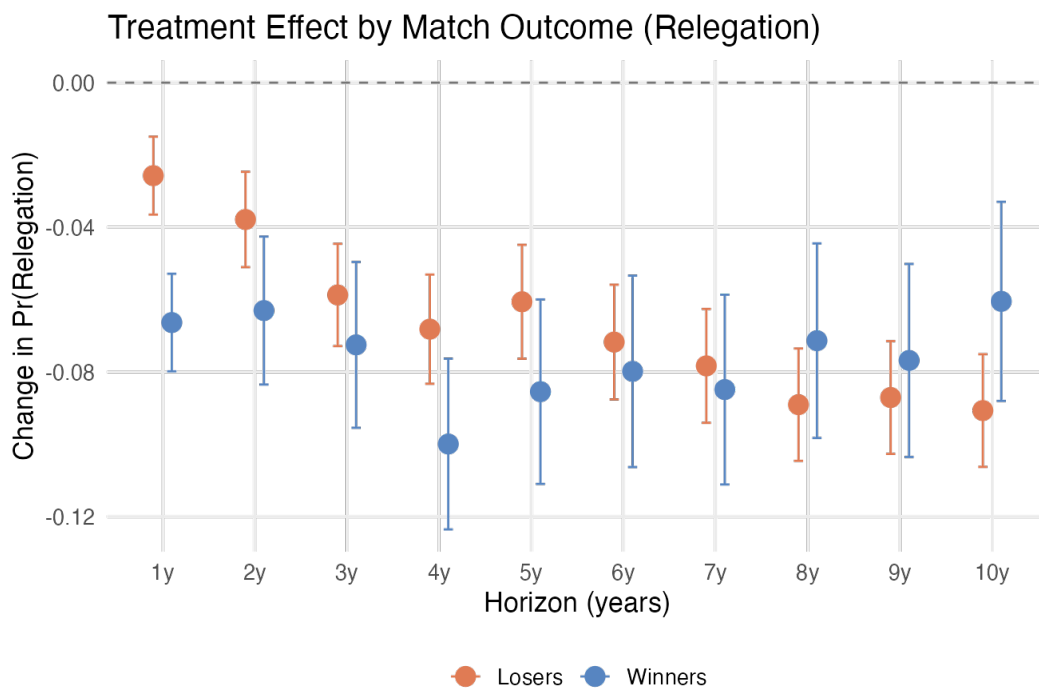


Figure 6: Treatment effect on relegation, decomposed by match outcome. Blue points represent clubs that won their FA Cup match; red points represent clubs that lost. Losing clubs receive the financial windfall without the confounding psychological effects of an upset victory. Negative values indicate protection against relegation.

Table 4 presents the full set of regression estimates underlying Figure 6.

7.2 Organizational Capacity as a Moderator.

Does the ability to translate a financial windfall into improved performance depend on a club’s pre-existing organizational capacity? As discussed in Section 2, a general principle from the literatures on absorptive capacity, dynamic capabilities, and resource windfalls is that organizations can only deploy new inflows productively when they possess sufficient complementary assets—personnel, infrastructure, and strategic know-how. Without such capacity, windfalls may be wasted or dissipated.

We test this by interacting the treatment indicator with a proxy for organizational capacity: the

Table 4: Win/loss decomposition: Effect of windfall on relegation

Horizon	Winners			Losers		
	Coef.	SE	p	Coef.	SE	p
1 year	-0.066***	(0.013)	<0.001	-0.026*	(0.011)	0.017
2 years	-0.063**	(0.020)	0.002	-0.038**	(0.013)	0.004
3 years	-0.073**	(0.023)	0.002	-0.059***	(0.014)	<0.001
4 years	-0.100***	(0.024)	<0.001	-0.068***	(0.015)	<0.001
5 years	-0.085***	(0.026)	<0.001	-0.061***	(0.016)	<0.001
6 years	-0.080**	(0.026)	0.003	-0.072***	(0.016)	<0.001
7 years	-0.085**	(0.026)	0.001	-0.078***	(0.016)	<0.001
8 years	-0.071**	(0.027)	0.008	-0.089***	(0.016)	<0.001
9 years	-0.077**	(0.027)	0.004	-0.087***	(0.016)	<0.001
10 years	-0.060*	(0.028)	0.028	-0.091***	(0.016)	<0.001
N	5,293–5,171			5,293–5,171		

Notes: Each row reports the coefficient on the treatment indicator from a separate linear probability model of relegation at horizon k . Winners: treatment is an indicator for treated clubs that won their FA Cup match. Losers: treatment is an indicator for treated clubs that lost. Negative values indicate protection against relegation. Sample restricted to club-seasons with non-missing treatment status, matching the main analysis ($N = 5,349$); N falls modestly across horizons due to missing baseline division data. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

treated club’s own home stadium capacity, which reflects its pre-existing infrastructure. Clubs with larger stadiums tend to have more developed commercial operations, larger administrative staffs, better training facilities, and more sophisticated scouting networks. While imperfect, this measure captures meaningful variation in the organizational sophistication of lower-league clubs. The logic parallels Nohria and Gulati (1996), who argue that the relationship between organizational slack and innovation depends on the firm’s capacity to channel excess resources toward productive uses rather than wasteful ones.

We estimate the capacity interaction using places climbed as the dependent variable, which provides a continuous measure of upward mobility through the league pyramid. Figure 7 presents the results. The interaction between treatment and stadium capacity is positive at every horizon and statistically significant from two years onward (marginally so at one year, $p = 0.06$), with $p < 0.01$ from years 3–6 and $p < 0.001$ from year 8 onward. The interaction coefficient grows from 132 at two years to 457 at ten years (units: places climbed $\times 10^6$ per unit of stadium capacity), indicating that the moderating role of capacity becomes more important over time as organizations learn to deploy resources.

The returns to a windfall appear to depend on the organizational structures available to deploy it, and this advantage compounds over time. The result parallels findings at the country level, where Mehlum et al. (2006) show that the resource curse is concentrated in countries with weak institutions, while countries with strong institutions convert resource wealth into growth. It also provides causal evidence for the absorptive capacity principle of Cohen and Levinthal (1990) in a new domain: financial rather than knowledge-based inflows, in a setting with random treatment assignment.

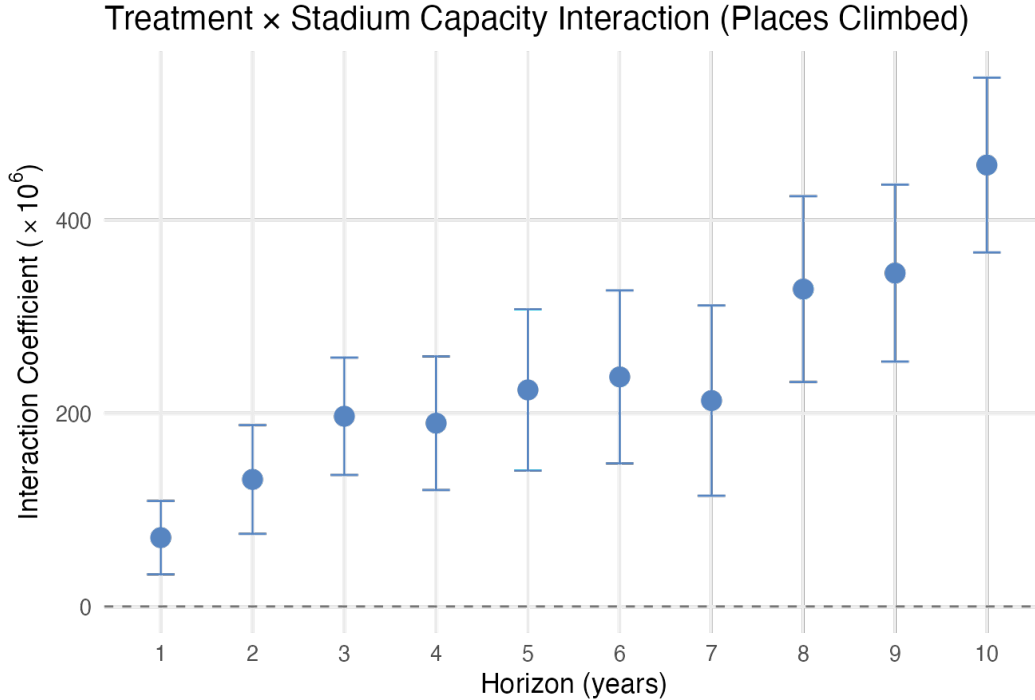


Figure 7: Moderating effect of organizational capacity (stadium capacity) on the treatment effect on places climbed. The interaction term captures whether clubs with greater pre-existing infrastructure benefit more from the financial windfall.

Capacity moderates upward mobility but not survival. We also estimate the capacity interaction using relegation as the dependent variable and find a markedly different pattern: the interaction is weak and statistically insignificant at most horizons, reaching significance only at the ten-year horizon ($p = 0.008$). The contrast suggests a distinction between a *floor effect* and a *ceiling effect*. Protection against relegation—preventing decline—may require only a financial buffer, which operates regardless of organizational capacity. But climbing the pyramid requires the organizational structures to identify and execute productive investments: better players, improved training facilities, stronger scouting networks. In our data, a financial buffer protects against decline regardless of organizational capacity, but moving up requires the structures to put the money to work. The capacity interaction with promotion on the Divisions 2+ subsample shows a similar but weaker pattern, reaching significance at scattered horizons (Appendix Table 9).

Robustness to alternative capacity measures. To assess whether the organizational capacity finding depends on the choice of proxy, we re-estimate the capacity interaction using three alternative measures that capture different dimensions of institutional development. First, *club age* (years since founding) captures the accumulation of organizational routines and institutional memory; younger organizations lack the established procedures and external legitimacy that develop over time (Stinchcombe 1965, Nelson and Winter 1982). Second, *historical peak division*—the highest tier of English football a club has ever reached—captures prior exposure to elite competition and

the organizational capabilities it requires, consistent with the idea that organizations with deeper institutional histories are better positioned to deploy new resources (Cohen and Levinthal 1990, Zahra and George 2002). Third, *division stability* (the standard deviation of a club’s league tier over the preceding five seasons) captures organizational reliability and consistency, which Hannan and Freeman (1984) argue signal underlying institutional quality.

All three alternative measures yield the same qualitative pattern as stadium capacity: the interaction with treatment is negative and strengthens at longer horizons, indicating that clubs with greater organizational depth benefit disproportionately from Cup windfalls (Appendix Tables 10–12). Historical peak division produces the strongest signal, with the interaction reaching statistical significance at the 0.1% level from six years onward. When we jointly estimate interactions with both the treated club’s own stadium capacity (organizational proxy) and the opponent’s stadium capacity (which determines windfall magnitude through gate receipts), the organizational capacity interaction is significant at every horizon while the windfall-dose interaction is weak (Appendix Table 13), consistent with the heterogeneity being driven by the club’s ability to deploy resources rather than by the amount received.

Table 5 presents the regression estimates for the capacity interaction with places climbed.

Table 5: Organizational capacity interaction: Treatment \times stadium capacity (places climbed)

Horizon	Coef. ($\times 10^6$)	SE ($\times 10^6$)	p	N
1 year	71.3	(38.0)	0.061	4,659
2 years	132.0*	(56.3)	0.020	4,553
3 years	197.0**	(60.7)	0.001	4,459
4 years	190.0**	(69.2)	0.006	4,363
5 years	224.0**	(83.4)	0.007	4,274
6 years	238.0**	(89.5)	0.008	4,198
7 years	213.0*	(98.5)	0.031	4,104
8 years	329.0***	(96.2)	<0.001	4,016
9 years	345.0***	(91.6)	<0.001	3,895
10 years	457.0***	(90.5)	<0.001	3,812

Notes: Each row reports the coefficient on the interaction term (BinTreat \times treated club’s stadium capacity) from a regression of places climbed at horizon k . Coefficients are scaled by 10^6 for readability; for example, a coefficient of 203 implies that a 10,000-seat increase in the treated club’s home stadium capacity is associated with 2.03 additional places climbed. Sample restricted to Divisions 1–5 (stadium capacity data unavailable for non-league clubs). N varies by horizon because places climbed requires non-missing rank data. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

8 Discussion and Conclusion

Does money improve organizational performance? Despite its centrality to management, economics, and organizational theory, this question has proven difficult to answer causally because financial resources and performance are jointly determined. By exploiting the random draw of the FA Cup—a natural experiment repeated thousands of times over seven decades—we provide credible causal evidence that financial windfalls improve the league position of small football clubs. Treated clubs climb approximately 18 league places over ten years, are significantly less likely to be relegated, and—among clubs for which promotion is possible—are more likely to be promoted. The effects grow monotonically over the ten-year horizon, consistent with reinvestment and compounding rather than a one-time boost.

First, the decomposition by match outcome provides direct evidence that the mechanism is financial rather than psychological. Clubs that lose their FA Cup match are significantly protected against relegation, with effects growing from 2 to nearly 9 percentage points over the decade—a pattern inconsistent with a psychological mechanism, which would dissipate rather than compound. At longer horizons, losers show larger point estimates than winners, and the year-1 gap between the two groups likely reflects additional financial treatment from advancing in the cup rather than pure momentum. Because losing clubs receive money without any confounding boost from victory, this decomposition addresses a concern that pervades the campaign finance literature (Jacobson 1978, Levitt 1994)—that observed correlations between resources and success may reflect unobserved quality—in a way that observational studies cannot. A falsification test provides further support: we estimate the same specification on the opponent club’s subsequent league position—the top-division club that hosted the match but received no meaningful financial shock from it. The effect is null at every horizon (Appendix Table 16).

Second, the moderating role of organizational capacity extends to financial resources a principle that Cohen and Levinthal (1990) established for knowledge: the returns to new inflows depend on pre-existing complementary assets. But our evidence goes further than the absorptive capacity literature by revealing an asymmetry. Capacity moderates upward mobility—clubs with stronger infrastructure benefit disproportionately—but not protection against decline, which does not appear to depend on organizational capacity. This floor-versus-ceiling distinction helps reconcile seemingly contradictory results in the windfall literature: the waste documented by Blanchard et al. (1994) and the dissipation found by Brollo et al. (2013) and Caselli and Michaels (2013) may reflect the absence of organizational capacity needed for *advancement*, even as the same windfalls provide a buffer against decline. The contrast with individual-level evidence sharpens the point: Carlson et al. (2015) find that NFL players go bankrupt at rates unrelated to the size of their career earnings, suggesting that without organizational structures to channel resources, even large income spikes fail to produce lasting financial health.

This interpretation also addresses the organizational slack debate (Cyert and March 1963, Nohria and Gulati 1996). Nohria and Gulati (1996) propose that the relationship between slack and innovation is curvilinear—beneficial up to a point, then harmful. Our heterogeneous treatment effects

are consistent with this view: financial windfalls improve league position on average, but the benefit for upward mobility is concentrated among clubs with greater organizational capacity, while protection against decline is unconditional. For the resource-based view of the firm (Barney 1991), our findings address a prior question: before asking whether resources confer *sustained* competitive advantage, we show that exogenous increases in financial resources can improve performance—but that the nature of the improvement depends on organizational capacity.

These results also speak to a broader puzzle. The windfall literature has produced mostly negative findings—waste in firms (Blanchard et al. 1994), corruption in governments (Brollo et al. 2013, Caselli and Michaels 2013), bankruptcy among individuals (Carlson et al. 2015)—which, taken together, might suggest that money simply does not improve performance. But the settings that produce negative results share a common feature: they lack one or more of the conditions under which money can be put to productive use. Football clubs have a clear, measurable objective (improve league position), face competitive pressure that punishes waste (a club that squanders a windfall loses matches and risks relegation), and possess organizational structures (managers, coaching staff, scouting networks, facilities) capable of turning money into better performance. The firms in Blanchard et al. (1994) faced weak competitive discipline over how to spend lawsuit proceeds. The governments in Brollo et al. (2013) and Caselli and Michaels (2013) lacked accountability, allowing revenues to fund corruption rather than public goods. The NFL players in Carlson et al. (2015) had no organizational structure at all. Evidence from public education points the same way: Jackson et al. (2016) find that court-ordered increases in school funding improve students’ long-run outcomes—again, a setting where organizations with clear missions and accountability turn resources into results. The upshot is not that windfalls are generally beneficial or generally wasteful, but that the effect of money on performance depends on whether the receiving organization has the objectives, discipline, and structure to use it well.

Third, the setting itself offers advantages that are difficult to replicate in studies of countries, governments, or large firms. Treatment is genuinely random, outcomes are unambiguous, the organizations share a common objective, and the sample is large enough to detect economically meaningful effects with reasonable precision.

Our study has limitations. We do not observe how clubs spend their windfall revenue, and so we cannot directly identify the investment channel through which money translates into improved league position. This is a deliberate feature of the reduced-form approach rather than a limitation alone: the productive use of a windfall is era-specific, varying with the labour market, regulatory environment, and competitive landscape. A club in the 1960s might sign two players with its windfall; a club in the 1990s might fund the conversion of a standing terrace to meet all-seater requirements; a club in the 2010s might hire analytics staff, invest in sports science and nutrition programmes, or cover a few months of a single player’s wages. Tracing spending channels would require era-specific models and would sacrifice the generality of the reduced-form estimate, which captures the effect of money across all these regimes. Stadium capacity is a coarse proxy for organizational capacity, and future work could incorporate richer measures of club governance,

staffing, and financial management. Furthermore, while the FA Cup draw is random, the set of clubs that reach the third round proper is not, though this selection operates identically for treatment and control clubs.

These limitations suggest directions for further work. Era-specific analyses could examine whether the composition of post-windfall spending has changed as English football's economics have evolved, whether the quality of club leadership moderates the effect, and whether there are circumstances under which windfalls are counterproductive—for example, by creating internal conflicts over how to allocate unexpected resources. The FA Cup setting, with its combination of random treatment, clear outcomes, and long time series, provides a useful setting for studying these questions.

Our findings suggest that the answer to this question depends on what kind of help is needed. A financial buffer appears sufficient to protect against decline, but moving up requires the organizational capacity to put money to use. The interaction between financial resources, organizational capacity, and the nature of the performance objective deserves further attention.

References

- Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management* 17(1), 99–120.
- Berset, S. and M. Schelker (2020). Fiscal windfall curse. *European Economic Review* 130, 103592.
- Blanchard, O. J., F. López-de Silanes, and A. Shleifer (1994). What do firms do with cash windfalls? *Journal of Financial Economics* 36(3), 337–360.
- Brollo, F., T. Nannicini, R. Perotti, and G. Tabellini (2013). The political resource curse. *American Economic Review* 103(5), 1759–1796.
- Cameron, A. C. and D. L. Miller (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources* 50(2), 317–372.
- Carlson, K., J. Kim, A. Lusardi, and C. F. Camerer (2015). Bankruptcy rates among NFL players with short-lived income spikes. *American Economic Review* 105(5), 381–386.
- Caselli, F. and G. Michaels (2013). Do oil windfalls improve living standards? Evidence from Brazil. *American Economic Journal: Applied Economics* 5(1), 208–238.
- Cohen, W. M. and D. A. Levinthal (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly* 35(1), 128–152.
- Cyert, R. M. and J. G. March (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice Hall.

- Dobson, S. and J. Goddard (2001). *The Economics of Football*. Cambridge: Cambridge University Press.
- Gerber, A. S. (1998). Estimating the effect of campaign spending on senate election outcomes using instrumental variables. *American Political Science Review* 92(2), 401–411.
- Hannan, M. T. and J. Freeman (1984). Structural inertia and organizational change. *American Sociological Review* 49(2), 149–164.
- Jackson, C. K., R. C. Johnson, and C. Persico (2016). The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. *Quarterly Journal of Economics* 131(1), 157–218.
- Jacobson, G. C. (1978). The effects of campaign spending in congressional elections. *American Political Science Review* 72(2), 469–491.
- Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review* 76(2), 323–329.
- Kleven, H. J., C. Landais, and E. Saez (2013). Taxation and international migration of superstars: Evidence from the European football market. *American Economic Review* 103(5), 1892–1924.
- Levitt, S. D. (1994). Using repeat challengers to estimate the effect of campaign spending on election outcomes in the U.S. House. *Journal of Political Economy* 102(4), 777–798.
- Mehlum, H., K. Moene, and R. Torvik (2006). Institutions and the resource curse. *The Economic Journal* 116(508), 1–20.
- Nelson, R. R. and S. G. Winter (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.
- Nohria, N. and R. Gulati (1996). Is slack good or bad for innovation? *Academy of Management Journal* 39(5), 1245–1264.
- Popplewell, M. J. (1986). *Committee of Inquiry into Crowd Safety and Control at Sports Grounds: Final Report*. London: HMSO. Cmnd 9710.
- Robinson, J. A., R. Torvik, and T. Verdier (2006). Political foundations of the resource curse. *Journal of Development Economics* 79(2), 447–468.
- Ross, M. L. (2001). Does oil hinder democracy? *World Politics* 53(3), 325–361.
- Sachs, J. D. and A. M. Warner (1995). Natural resource abundance and economic growth. *NBER Working Paper* (5398).
- Stinchcombe, A. L. (1965). Social structure and organizations. In J. G. March (Ed.), *Handbook of Organizations*, pp. 142–193. Chicago: Rand McNally.

- Szymanski, S. (2000). A market test for discrimination in the English professional soccer leagues. *Journal of Political Economy* 108(3), 590–603.
- Szymanski, S. (2003). The economic design of sporting contests. *Journal of Economic Literature* 41(4), 1137–1187.
- Szymanski, S. and T. Kuypers (1999). *Winners and Losers: The Business Strategy of Football*. London: Viking.
- Szymanski, S. and R. Smith (1997). The English football industry: Profit, performance and industrial structure. *International Review of Applied Economics* 11(1), 135–153.
- Taylor, L. J. (1990). *The Hillsborough Stadium Disaster: Final Report*. London: HMSO. Cm 962.
- Teece, D. J., G. Pisano, and A. Shuen (1997). Dynamic capabilities and strategic management. *Strategic Management Journal* 18(7), 509–533.
- UK Parliament (1975). Safety of sports grounds act 1975. c. 52.
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48(4), 817–838.
- Zahra, S. A. and G. George (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review* 27(2), 185–203.

A Pre-Treatment Trajectory Analysis

We merge historical league standings data with our FA Cup sample to construct measures of each club’s trajectory through the football pyramid in the three to five seasons preceding the FA Cup draw. For each club-season observation, we obtain the club’s tier and within-tier position for each of the five preceding seasons, yielding up to six data points per observation (the season of the draw plus five lookback years). Merge coverage is approximately 93–95 percent, with missing observations concentrated among non-league clubs below the fourth tier. We construct the following trajectory variables:

- *Division change (3-year, 5-year)*: The difference in tier between three or five years before the draw and the season of the draw. Positive values indicate upward movement (the club occupied a lower tier in the past).
- *Was promoted / relegated (3-year, 5-year)*: Binary indicators for whether the club experienced any year-over-year tier change (up or down) within the window.
- *Places climbed (3-year, 5-year)*: Change in global league position (sum of teams in higher divisions plus within-division rank) over the window. Positive values indicate improvement.
- *Trajectory slope (5-year)*: The negative of the OLS slope of tier on time over the six pre-treatment data points. Positive values indicate an improving trajectory.

Table 6: Balance on pre-treatment trajectory variables

Variable	Unconditional			Conditional on Division			<i>N</i>
	Diff.	SE	<i>p</i>	Diff.	SE	<i>p</i>	
Div. change (3yr)	−0.067	(0.031)	0.034	0.004	(0.031)	0.910	2,819
Div. change (5yr)	−0.109	(0.037)	0.003	0.005	(0.036)	0.895	2,781
Was promoted (3yr)	−0.068	(0.018)	<0.001	−0.022	(0.018)	0.229	3,026
Was promoted (5yr)	−0.082	(0.020)	<0.001	−0.016	(0.020)	0.424	3,026
Was relegated (3yr)	−0.024	(0.019)	0.212	−0.018	(0.020)	0.349	3,026
Was relegated (5yr)	−0.044	(0.021)	0.032	−0.027	(0.021)	0.195	3,026
Places climbed (3yr)	−0.567	(0.664)	0.394	0.232	(0.666)	0.727	2,819
Places climbed (5yr)	−1.489	(0.819)	0.069	0.371	(0.808)	0.647	2,781
Trajectory slope (5yr)	−0.019	(0.008)	0.010	0.002	(0.007)	0.750	2,857

Notes: Each row reports the coefficient on the binary treatment indicator from a regression of the trajectory variable on treatment status. “Unconditional” columns report results from a bivariate regression; “Conditional on Division” columns add division fixed effects. The sample consists of clubs in Divisions 2–4 with non-missing standings data for the relevant lookback window. Heteroskedasticity-consistent (HC1) standard errors in parentheses.

Table 6 presents balance tests for all nine trajectory variables. In unconditional comparisons, treated clubs exhibit significantly worse pre-treatment trajectories on several dimensions: they were less likely to have been recently promoted and show a declining trajectory slope. These differences are entirely compositional. Because treatment probability is higher for lower-division clubs—who

have more potential higher-division opponents in the draw—the treated group overrepresents clubs from Divisions 3 and 4, which have mechanically different trajectory distributions than Division 2 clubs. Conditional on division, every imbalance vanishes: the largest t -statistic is 1.28 and no p -value falls below 0.19. The draw is random, and treated and control clubs within the same division are on indistinguishable pre-treatment trajectories.

Table 7: Pre-treatment trajectory as a moderator of the treatment effect

Moderator		Horizon			
		3 years	5 years	7 years	10 years
Div. change (3yr)	$\hat{\gamma}$	−0.019	−0.030	−0.001	0.031
	SE	(0.026)	(0.025)	(0.026)	(0.025)
	p	0.447	0.232	0.976	0.224
Div. change (5yr)	$\hat{\gamma}$	−0.036	−0.037	−0.017	−0.001
	SE	(0.022)	(0.021)	(0.022)	(0.022)
	p	0.104	0.088	0.454	0.952
Trajectory slope (5yr)	$\hat{\gamma}$	−0.188	−0.193	−0.066	−0.046
	SE	(0.103)	(0.099)	(0.104)	(0.099)
	p	0.068	0.051	0.522	0.641
Was promoted (3yr)	$\hat{\gamma}$	−0.013	−0.029	−0.010	0.034
	SE	(0.037)	(0.038)	(0.038)	(0.039)
	p	0.726	0.435	0.791	0.391
Was relegated (3yr)	$\hat{\gamma}$	0.016	−0.009	−0.031	−0.050
	SE	(0.039)	(0.040)	(0.040)	(0.040)
	p	0.692	0.833	0.442	0.214

Notes: Each cell reports the coefficient $\hat{\gamma}$ on the interaction term $\text{BinTreat} \times \text{Moderator}$ from the model $\text{Promoted}_{K_y} = \alpha + \beta \cdot \text{BinTreat} + \delta \cdot \text{Moderator} + \gamma \cdot \text{BinTreat} \times \text{Moderator} + \text{Division FE} + \varepsilon$. The outcome is cumulative promotion within K years. Division fixed effects are included to isolate trajectory from division composition. The sample consists of clubs in Divisions 2–4 with non-missing trajectory data. Heteroskedasticity-consistent (HC1) standard errors in parentheses. No interaction reaches conventional significance ($p < 0.05$) at any horizon.

Table 7 tests whether pre-treatment trajectory moderates the treatment effect. The interaction between treatment and each trajectory variable is insignificant at every horizon for all five moderators. The five-year division change and trajectory slope show a suggestive negative pattern at intermediate horizons (3–5 years), hinting that clubs already on an upward trajectory may benefit marginally less from the windfall, but no coefficient reaches conventional significance. The overall picture is clear: rising and declining clubs benefit similarly from windfalls, and the treatment effect does not depend on pre-treatment momentum. Combined with the balance results, this analysis validates the identification strategy and confirms that the unconditional treatment estimate is not biased by compositional differences in pre-treatment trajectories.

Table 8: Win/loss decomposition: Effect of windfall on places climbed

Horizon	Winners			Losers		
	Coef.	SE	p	Coef.	SE	p
1 year	10.0***	(2.28)	<0.001	2.8***	(0.80)	<0.001
2 years	9.3***	(2.48)	<0.001	3.7***	(1.04)	<0.001
3 years	5.8*	(2.48)	0.019	6.2***	(1.15)	<0.001
4 years	12.5***	(2.80)	<0.001	8.4***	(1.36)	<0.001
5 years	12.4***	(3.16)	<0.001	10.0***	(1.55)	<0.001
6 years	15.6***	(3.64)	<0.001	10.1***	(1.60)	<0.001
7 years	18.7***	(3.89)	<0.001	11.7***	(1.77)	<0.001
8 years	20.3***	(3.92)	<0.001	12.9***	(1.73)	<0.001
9 years	21.7***	(4.16)	<0.001	13.6***	(1.76)	<0.001
10 years	21.6***	(4.23)	<0.001	14.4***	(1.79)	<0.001

Notes: Each row reports the coefficient on the treatment indicator from a separate regression of places climbed (net change in league position) at horizon k . Winners: treatment is an indicator for treated clubs that won their FA Cup match. Losers: treatment is an indicator for treated clubs that lost. Positive values indicate upward movement through the pyramid. Sample restricted to club-seasons with non-missing treatment status, matching the main analysis; N ranges from 4,111 (10-year) to 5,013 (1-year) due to missing rank data at longer horizons. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

B Robustness to Club-Clustered Standard Errors

Our main specification treats each FA Cup draw as an independent randomization, which motivates heteroskedasticity-consistent (HC1) standard errors. However, because the same club can appear in multiple draws across our 70-year sample—186 unique clubs appear with a median of 15 and a mean of 29 appearances each—outcomes for the same club may be serially correlated. A club’s league trajectory reflects persistent organizational characteristics that generate within-club correlation regardless of FA Cup outcomes. At longer horizons this is especially relevant: a club’s position ten years after a 1960 draw and ten years after a 1965 draw are partly measuring the same underlying trajectory.

To address this, we re-estimate all main specifications clustering standard errors at the club level (Cameron and Miller 2015). Tables 14 and 15 report results. Clustered standard errors are moderately larger at long horizons (the SE ratio reaches $1.5\times$ at 10 years for places climbed), reflecting within-club persistence, but all conclusions are unchanged: the treatment effect on places climbed and relegation avoidance remains significant at every horizon ($p < 0.001$ for all places-climbed specifications; $p < 0.001$ at all horizons for relegation), and the organizational capacity interaction remains significant from year 1 onward.

Table 9: Organizational capacity interaction: Treatment \times stadium capacity (cumulative promotion, Divisions 2+ sample)

Horizon	Coef. ($\times 10^6$)	SE ($\times 10^6$)	p	N
1 year	1.3	(1.1)	0.218	3,144
2 years	2.2	(1.2)	0.068	3,144
3 years	1.8	(1.3)	0.163	3,144
4 years	3.0*	(1.3)	0.021	3,144
5 years	1.5	(1.3)	0.248	3,144
6 years	2.5*	(1.3)	0.049	3,144
7 years	2.5	(1.3)	0.053	3,144
8 years	2.4	(1.3)	0.066	3,144
9 years	2.5	(1.3)	0.054	3,144
10 years	3.7**	(1.3)	0.004	3,144

Notes: Each row reports the coefficient on the interaction term (BinTreat \times treated club's stadium capacity) from a linear probability model of cumulative promotion at horizon k . The sample is restricted to Divisions 2+ (promotion is mechanically impossible for top-division clubs) and to Divisions 2–5 (stadium capacity data unavailable for non-league clubs). Coefficients are scaled by 10^6 for readability. The interaction is significant at scattered horizons (4, 6, and 10 years), consistent with but weaker than the places-climbed results in Table 5. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 10: Alternative capacity measure: Treatment \times club age (places climbed)

Horizon	Coef. (interaction)	SE	p	N
1 year	-0.015	(0.025)	0.547	4,985
2 years	-0.045	(0.035)	0.194	4,873
3 years	-0.114**	(0.039)	0.003	4,773
4 years	-0.155***	(0.045)	<0.001	4,676
5 years	-0.145**	(0.049)	0.003	4,582
6 years	-0.083	(0.050)	0.098	4,501
7 years	-0.125*	(0.057)	0.028	4,399
8 years	-0.140**	(0.054)	0.010	4,303
9 years	-0.173**	(0.057)	0.002	4,178
10 years	-0.193**	(0.061)	0.002	4,095

Notes: Each row reports the coefficient on the interaction term (BinTreat \times club age) from a regression of places climbed at horizon k . Club age is measured as the year of the FA Cup match minus the club's founding year. A negative interaction indicates that older clubs benefit more from the windfall. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 11: Alternative capacity measure: Treatment \times historical peak division (places climbed)

Horizon	Coef. (interaction)	SE	p	N
1 year	-1.72	(1.01)	0.090	4,838
2 years	-3.55**	(1.35)	0.009	4,730
3 years	-3.62*	(1.42)	0.011	4,634
4 years	-4.22**	(1.58)	0.007	4,540
5 years	-5.26**	(1.78)	0.003	4,451
6 years	-6.55***	(1.84)	<0.001	4,374
7 years	-9.25***	(2.04)	<0.001	4,276
8 years	-8.76***	(2.00)	<0.001	4,184
9 years	-8.90***	(2.09)	<0.001	4,064
10 years	-11.19***	(2.13)	<0.001	4,005

Notes: Each row reports the coefficient on the interaction term (BinTreat \times historical peak division) from a regression of places climbed at horizon k . Historical peak division is the highest tier (lowest division number) the club had reached in any prior season. Since lower numbers indicate higher divisions, a negative interaction means that clubs with historically stronger records benefit more from the windfall. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 12: Alternative capacity measure: Treatment \times division stability (places climbed)

Horizon	Coef. (interaction)	SE	p	N
1 year	0.14	(3.69)	0.969	4,740
2 years	0.16	(4.24)	0.971	4,635
3 years	-8.04	(4.80)	0.094	4,542
4 years	-10.42	(5.34)	0.051	4,451
5 years	-11.52	(6.12)	0.060	4,365
6 years	-12.92*	(6.28)	0.040	4,291
7 years	-16.61*	(7.07)	0.019	4,197
8 years	-20.23**	(6.89)	0.003	4,108
9 years	-22.93**	(7.18)	0.001	3,991
10 years	-23.38**	(7.36)	0.001	3,948

Notes: Each row reports the coefficient on the interaction term (BinTreat \times division stability) from a regression of places climbed at horizon k . Division stability is the standard deviation of the club's league tier over the five seasons preceding the FA Cup match; higher values indicate greater volatility. A negative interaction means that more stable clubs (lower SD) benefit more from the windfall. Requires five consecutive prior seasons of data. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 13: Dual capacity interaction: Treatment \times organizational capacity and Treatment \times windfall dose (places climbed)

Horizon	Org. capacity (own stadium)		Windfall dose (opponent stadium)		N
	Coef. ($\times 10^6$)	p	Coef. ($\times 10^6$)	p	
1 year	90.5**	0.007	-5.1	0.880	4,658
2 years	198.1***	<0.001	14.8	0.706	4,552
3 years	199.7***	<0.001	53.9	0.181	4,458
4 years	198.7***	<0.001	64.2	0.156	4,362
5 years	226.4***	<0.001	16.9	0.725	4,273
6 years	295.6***	<0.001	-18.8	0.703	4,197
7 years	324.8***	<0.001	-64.4	0.237	4,103
8 years	349.8***	<0.001	-146.3*	0.013	4,015
9 years	329.1***	<0.001	-101.3	0.077	3,894
10 years	389.7***	<0.001	-155.8*	0.012	3,811

Notes: Each row reports coefficients from a single regression of places climbed at horizon k that includes two interaction terms simultaneously: BinTreat \times treated club's own stadium capacity (organizational capacity) and BinTreat \times opponent's stadium capacity (windfall dose, since gate receipts from the FA Cup match are determined by the opponent's ground). The organizational capacity interaction is significant at every horizon, while the windfall-dose interaction is largely insignificant, indicating that the heterogeneity in treatment effects is driven by the club's pre-existing infrastructure rather than by the size of the windfall received. Coefficients scaled by 10^6 . Sample restricted to Divisions 1-5. Heteroskedasticity-consistent (HC1) standard errors. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 14: Main results with club-clustered standard errors

Horizon	Places Climbed			Relegation			N
	Coef.	SE	p	Coef.	SE	p	
1 year	3.32	(0.77)	<0.001	-0.034	(0.010)	0.001	5,013
2 years	5.01	(0.89)	<0.001	-0.045	(0.012)	<0.001	4,891
3 years	7.91	(1.21)	<0.001	-0.067	(0.014)	<0.001	4,790
4 years	10.20	(1.40)	<0.001	-0.074	(0.017)	<0.001	4,684
5 years	11.62	(1.63)	<0.001	-0.066	(0.017)	<0.001	4,589
6 years	11.78	(1.80)	<0.001	-0.080	(0.018)	<0.001	4,508
7 years	13.98	(2.09)	<0.001	-0.088	(0.018)	<0.001	4,420
8 years	14.83	(2.06)	<0.001	-0.096	(0.018)	<0.001	4,325
9 years	16.70	(2.37)	<0.001	-0.103	(0.019)	<0.001	4,200
10 years	17.83	(2.61)	<0.001	-0.106	(0.018)	<0.001	4,111

Notes: Each row reports the coefficient on the binary treatment indicator from a separate linear probability model (relegation) or OLS regression (places climbed). Standard errors clustered by club (186 clusters) in parentheses. Compare to main Tables 1-2, which report HC1 standard errors. Point estimates are identical; only standard errors and p -values change.

Table 15: Organizational capacity interaction with club-clustered standard errors

Horizon	Coef. ($\times 10^6$)	SE ($\times 10^6$)	p	N
1 year	134	(46)	0.004	5,002
2 years	219	(55)	<0.001	4,881
3 years	317	(87)	<0.001	4,778
4 years	383	(118)	0.001	4,672
5 years	507	(156)	0.001	4,578
6 years	555	(170)	0.001	4,497
7 years	712	(209)	0.001	4,409
8 years	712	(206)	0.001	4,315
9 years	728	(228)	0.001	4,190
10 years	827	(255)	0.001	4,102

Notes: Each row reports the coefficient on the interaction term (BinTreat \times treated club's stadium capacity) from a regression of places climbed at horizon k . Standard errors clustered by club (186 clusters). Compare to main Table 5, which reports HC1 standard errors. All interaction terms remain significant, indicating that the organizational capacity moderation is robust to within-club serial correlation.

C Falsification Test: Opponent Club’s Outcomes

If the treatment effect on small clubs operates through the financial windfall from playing in a large stadium, then the *opponent* club—the top-division host—should show no effect, since it receives no meaningful differential revenue from the draw. We test this by estimating the same specification with the top-division club’s subsequent league position as the outcome.

Table 16: Falsification test: Effect on opponent (top-division) club’s league position

Horizon	Panel A: Places Climbed				Panel B: Relegation			
	Coef.	SE	<i>p</i>	<i>N</i>	Coef.	SE	<i>p</i>	<i>N</i>
1 year	0.894	(0.597)	0.134	1,635	−0.021	(0.018)	0.232	1,657
2 years	1.334	(0.891)	0.134	1,608	−0.022	(0.021)	0.299	1,657
3 years	1.581	(1.131)	0.162	1,581	−0.020	(0.023)	0.373	1,657
4 years	0.781	(1.389)	0.574	1,553	−0.002	(0.024)	0.950	1,657
5 years	0.783	(1.562)	0.616	1,530	0.002	(0.025)	0.944	1,657
6 years	2.114	(1.573)	0.179	1,505	−0.010	(0.025)	0.690	1,657
7 years	2.047	(1.629)	0.209	1,476	−0.022	(0.025)	0.374	1,657
8 years	1.869	(1.653)	0.258	1,449	−0.025	(0.025)	0.325	1,657
9 years	0.918	(1.734)	0.596	1,420	0.000	(0.025)	0.985	1,657
10 years	1.365	(1.857)	0.462	1,393	−0.006	(0.025)	0.825	1,657

Notes: Each row reports the coefficient on the binary treatment indicator from a separate regression, where the outcome is measured for the *opponent* (top-division) club rather than the treated (lower-division) club. Treatment is an indicator for hosting a lower-division opponent assigned by the FA Cup draw. These clubs receive no meaningful differential revenue from the match. Panel A: places climbed; Panel B: relegation. Sample restricted to Division 1 clubs. No coefficient reaches conventional significance at any horizon, consistent with the hypothesis that the treatment effect on small clubs operates through the financial windfall rather than through the match itself. Heteroskedasticity-consistent (HC1) standard errors in parentheses.

D Robustness to Repeated Treatment

Many clubs appear in the FA Cup draw multiple times across the sample period. To ensure that the main results are not confounded by prior treatments, we restrict the sample to each club’s first treatment and drop all subsequent observations for that club. This yields a smaller sample ($N \approx 900$) in which every treated club is receiving its first windfall and no control club has been previously treated.

Table 17: First-treatment restriction: Effect of windfall on league position

Horizon	Panel A: Places Climbed				Panel B: Relegation			
	Coef.	SE	p	N	Coef.	SE	p	N
1 year	5.9*	(2.45)	0.016	771	-0.039	(0.021)	0.066	897
2 years	7.6*	(3.36)	0.025	749	-0.079**	(0.027)	0.003	881
3 years	12.9***	(3.01)	<0.001	740	-0.121***	(0.029)	<0.001	870
4 years	17.7***	(3.62)	<0.001	731	-0.146***	(0.031)	<0.001	856
5 years	18.6***	(4.62)	<0.001	724	-0.121***	(0.035)	<0.001	853
6 years	19.0***	(4.85)	<0.001	723	-0.127***	(0.036)	<0.001	852
7 years	25.2***	(5.64)	<0.001	710	-0.171***	(0.034)	<0.001	846
8 years	30.4***	(4.78)	<0.001	703	-0.189***	(0.033)	<0.001	845
9 years	34.1***	(4.81)	<0.001	672	-0.214***	(0.032)	<0.001	849
10 years	38.1***	(5.11)	<0.001	669	-0.212***	(0.033)	<0.001	848

Notes: Each row reports the coefficient on the binary treatment indicator from a separate regression. The sample is restricted to each club’s first FA Cup treatment: for every club that is ever treated, all observations after its first treatment year are dropped, so that no treated club has received a prior windfall and no control observation is contaminated by earlier treatment. Panel A: places climbed; Panel B: relegation. Effects are approximately twice as large as the full-sample estimates (Table 2), consistent with repeated treatments attenuating the signal. Heteroskedasticity-consistent (HC1) standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

E Stadium Capacity Data

Our stadium capacity data assign time-varying capacities to 101 stadiums that experienced significant capacity changes during the sample period, yielding 309 stadium-period observations (median 3 periods per stadium, range 1–6). A further 115 smaller stadiums retain a single fixed capacity throughout. The stadiums with time-varying data account for approximately 86 percent of match observations.

Three regulatory interventions drove the largest capacity changes. First, the Safety of Sports Grounds Act (1975) introduced mandatory safety certificates for grounds holding more than 10,000 spectators, triggering substantial reductions in the late 1970s: Ninian Park’s capacity fell from approximately 50,000 to 20,000, and Goodison Park’s from 56,000 to 35,000. Second, the Bradford fire of 1985 and the subsequent Popplewell Report (Popplewell 1986) prompted further restrictions via the Fire Safety and Safety of Places of Sport Act 1987, with particularly severe effects on older wooden stands—Elm Park fell from 20,000 to 8,000, and Belle Vue from 25,000 to 7,000.

Third, the Taylor Report (Taylor 1990), issued after the 1989 Hillsborough disaster, mandated all-seater stadiums for the top two divisions by August 1994. Converting standing terraces to seated accommodation sharply reduced capacities—the Anfield Kop, for instance, went from roughly 30,000 standing to 12,390 seated—though some clubs later rebuilt stands and recovered capacity.

Stadium capacities and period boundaries were compiled from Wikipedia’s articles on individual grounds, supplemented by specialist databases (StadiumDB, Football Ground Guide) and club-specific historical sources. Per-stadium sources are recorded in the replication dataset.