Event Alpha and Event Risk Premia

Sofonias Alemu Korsaye
Johns Hopkins University, Carey Business School

Dacheng Xiu
University of Chicago Booth School of Business
National Bureau of Economic Research

August 13, 2026

Abstract

Returns around scheduled corporate and macroeconomic events are large and heterogeneous. We develop a characteristic-based latent-factor framework separating event-specific premia from compensation for common return variation across overlapping events. Pre-scheduled event exposure gives high-dimensional machine-learning factor models causal rather than predictive content. In U.S. stocks over 1980–2024, earnings announcements and FOMC decisions carry significant premia. Momentum, turnover, size, beta, and volatility explain substantial cross-sectional variation. Event scores estimated from prior data rank subsequent event-date returns. The framework unifies event studies, conditional factor models, and recurring-event portfolios.

Keywords: Scheduled news; event studies; macroeconomic announcements; corporate events; peer-announcement exposure; heterogeneous conditional returns.

JEL Classification: G12; G14; C21; C23; C32; C58.

1 Introduction

The link between systematic information events and asset prices is one of the oldest questions in financial economics. The seminal contribution of Fama et al. (1969) provided the framework that has guided five decades of empirical research: an event-time abnormal return that measures the component of an event firm's return not explained by a benchmark fitted outside the event window. Brown and Warner (1985) and MacKinlay (1997) formalized the inferential properties of this design, Kothari and Warner (2007) survey its econometrics, and the resulting toolkit has been applied to nearly every important news arrival in asset pricing and corporate finance. Earnings announcements, payout decisions, mergers, share repurchases, regulatory rulings, monetary-policy decisions, and macroeconomic releases are all studied through the returns concentrated around their event dates.

Two facts emerging from this literature motivate our analysis. First, scheduled events account for a disproportionate share of average stock returns. Narrow event windows contain economically large premia even though they represent only a small fraction of trading days. Second, event returns vary sharply across firms. Size, momentum, market exposure, liquidity, payout policy, and other characteristics predict which firms earn the largest returns around recurring events, and anomaly returns concentrate disproportionately on news and earnings days (Engelberg et al., 2018). An average abnormal return is therefore an important summary, but it leaves much of the economic content of the event unexplained.

The cross-section is central to distinguishing economic explanations for event returns. A premium shared uniformly by event firms resembles a calendar-wide reward for bearing event risk. A premium that rises with market beta or volatility points toward aggregate risk. A premium concentrated among small, inactive, or difficult-to-trade firms instead highlights information processing and limits to arbitrage. Payout characteristics may be especially important on distribution dates, while momentum and profitability may reveal which earnings events contain the most consequential information. These explanations make different predictions across firms even when they produce the same average event return. Recovering the conditional cross-section is therefore necessary for moving from the existence of an event premium to its economic organization.

This heterogeneity raises questions that the usual event study treats separately. Which firms earn the event premium? Does the ranking reflect an event-specific return or exposure to common factors? Does a firm's systematic exposure change when it enters an event window? How do firm events interact with market-wide announcements that occur on the same day? Classical event studies, pooled event regressions, and calendar-time portfolios provide compelling estimates of average abnormal returns. They do not jointly recover the cross-section of event premia, the common factor realization on the event date, and the way event exposure varies with firm characteristics.

We develop a framework that brings these objects into one model. Expected returns contain a baseline component, a firm-event component, and a market-wide announcement component. Firm characteristics determine how each component varies across stocks. A latent-factor structure captures common return movements and allows the loading on those movements to change when a firm experiences an event. Market-wide announcement premia are estimated separately as unrestricted vectors, so how much of each premium the factor-loading span carries is measurable. Coincidence observations—firm events that fall on macro-announcement days—receive their own flexible controls so that they do not distort the clean event comparisons. The traditional scalar event premium becomes a conditional cross-section: the framework measures who earns the premium, how much is specific to the event itself, and how much varies with common return conditions.

This decomposition gives content to the distinction between event alpha and event risk premia. Event alpha is the event-level component; the event risk-premium component is the part associated with the change in exposure to latent common factors. The standard regressed-PCA normalization of Chen et al. (2025) (hereafter CRW) identifies the two components, while their sum—the fitted

event return—does not depend on how the latent factors are centered. The decomposition therefore preserves the familiar event-study object while showing whether its cross-sectional variation is concentrated in the event level or in exposure to common factors.

The framework is built for recurring events rather than a single, isolated experiment. Earnings announcements, distribution dates, monetary-policy meetings, and macroeconomic data releases can be studied together while retaining their distinct economic roles. Firm and macro events may overlap, and multiple firm events may occur on the same date. The model uses these overlaps rather than discarding them: cross-sectional comparisons identify firm-event differences, while calendar variation identifies market-wide announcement premia. Peer, industry, and network exposures fit naturally into the same structure when spillovers are the economic object of interest.

Estimation mirrors this economic logic. Firms without a firm-specific event reveal the daily relation between characteristics and returns. Variation across announcement and nonannouncement dates isolates the market-wide component. The remaining common movements recover the latent factors. Event firms then reveal how both the level of returns and their factor exposures change during firm events. Each stage is transparent and scalable to a large unbalanced panel. The accompanying theory establishes identification, consistency, joint asymptotic normality, and bootstrap inference that carries uncertainty from the estimated benchmark and latent factors into the event-premium estimates.

The paper contributes in three ways. First, it generalizes the central event-study estimand from one average abnormal return to a conditional cross-section of event returns. This extension directly identifies the firm characteristics associated with economically large event premia. Second, it connects event studies to characteristic-based latent-factor models. The resulting decomposition distinguishes an event-level return from event-related variation in common exposure without requiring observed factors to span the relevant common return variation. Third, it converts the estimated cross-section into event scores. These scores provide a direct test of whether the estimated heterogeneity travels across firms and through time, and they connect the econometric decomposition to the calendar-time portfolios used throughout the event-study literature.

Conditional latent-factor models estimated from high-dimensional characteristic panels are the workhorse of machine-learning asset pricing (Kelly et al., 2019; Gu et al., 2021), and their output is evaluated predictively, by how well fitted conditional moments forecast returns out of sample. We estimate the same class of high-dimensional conditional models but ask a causal question: because event exposure is pre-scheduled and therefore predetermined, the latent-factor benchmark identifies event-specific premia separately from announcement-day changes in systematic exposure. The event scores submit the estimated cross-section to the same test the prediction literature applies: fitted without the evaluation firms or years, they rank realized returns on subsequent event dates.

We apply the framework to 42.3 million U.S. stock-day observations from 1980 through 2024. Firm events consist of earnings announcements and distribution ex-dates. Market-wide events are CPI, PPI, Employment Situation, and FOMC announcements. The event calendar is rich: firm

events provide within-day comparisons between event and nonevent stocks, while macro announcements provide common date-level variation. Joint estimation preserves the identity of each event rather than forcing earnings, distributions, and macro releases to share a common premium.

The combined calendar also creates comparisons that separate the roles of firm and aggregate news. On an ordinary earnings date, the model compares the event firm with otherwise similar nonevent firms exposed to the same market-wide return conditions. When earnings coincide with an FOMC decision or an employment release, the market-wide component is estimated from the cross-section and does not have to be absorbed into the earnings premium. Conversely, the macro coefficients use thousands of firms on the same announcement date while allowing their responses to differ with characteristics. This joint use of within-date and across-date variation is the empirical advantage of studying all recurring events in one system.

The estimated average premia are economically large. For a firm at the cross-sectional mean of every characteristic, the estimated premium is about 9.5 basis points around earnings announcements, 33.2 basis points on cash-distribution ex-dates, and 17.0 basis points on FOMC dates. CPI and PPI premia are close to zero, showing why disaggregation matters. The broad distribution category produces a 52-basis-point estimate because it includes a small set of high-return noncash records, whereas the cash-only four-stage estimate is 33 basis points.

The cross-section is at least as important as these averages. Firm-event premia vary most strongly with long-term momentum, turnover, payout characteristics, and size. Macro announcement premia vary with market beta, idiosyncratic volatility, size, and momentum. Earnings and distribution events have distinct characteristic profiles, and the FOMC component is concentrated among firms with high market exposure. The estimated macro premium carries a component outside the span of the estimated factor loadings—four fifths of its squared length lies in that span, one fifth does not—so we read the macro coefficients as reduced-form conditional-return premia rather than as announcement-day shifts in expected factor returns. Maintaining the factor-channel restriction instead, the minimum-distance estimate of the factor announcement premium is economically large on FOMC dates: 17% of the first factor's daily standard deviation (bootstrap $t = 2.9$). The third contrast of the two-way design, the firm-event effect on an announcement day, requires no such restriction: firm events command a positive and significant premium on every macro announcement day, 8.8 to 19.7 basis points at the average characteristic, carried by the announcement-day event level rather than by event-day exposure. Set against the clean-date effect within the same bootstrap draws, firm-event compensation is unchanged on CPI days but roughly halved on PPI and Employment days—the sign investor distraction predicts, rather than compounding announcement risk—and modestly lower on FOMC days. The latent-factor decomposition shows that the average firm-event premium is concentrated in the event-level component, while the factor-dependent component varies meaningfully across firms and dates. An extension to peer-announcement exposure finds that industry peers have an average 0.57-basis-point higher conditional return on the trading day after a bellwether earnings announcement ($t = 1.72$ under the operative industry-clustered inference; $t = 2.08$ under the baseline bootstrap), with the largest positive characteristic slopes on

profitability measures.

These estimates also have strong predictive content. In cross-sectional holdouts, event scores rank returns among firms that were not used to estimate the score. In an expanding-window exercise, parameters estimated from prior years continue to order subsequent event returns. High-score event portfolios outperform low-score portfolios, with an annualized Sharpe ratio of about 2.4 for the daily high-minus-low spread (Newey–West $t = 13.5$); aggregating that daily series to weekly and monthly frequencies leaves the ordering intact. Following the recurring-event literature, these portfolios summarize gross conditional event returns. Their ordering shows that event premia are not merely average calendar effects: they contain a persistent and economically important cross-sectional structure.

Taken together, the results show that scheduled events reorganize the cross-section of expected returns. The same calendar date can carry a market-wide premium, a firm-specific event premium, and a change in systematic exposure, with each component varying across firms. Measuring these components jointly provides a more complete account of where event returns arise and which firms earn them.

The paper builds on several distinct strands of literature.

The classical event-study apparatus has been refined to address two long-standing challenges that motivate the design choices in our framework. Calendar-time clustering of events contaminates per-firm pre-event windows when firm events overlap with one another's estimation periods; the multivariate regression model of Schipper and Thompson (1983) and Binder (1985), together with the conditioning extension of Thompson (1985), addresses this by pooling event firms in a single panel regression. Non-normality and serial dependence of returns motivate the calendar-time portfolio approach of Jaffe (1974), Mandelker (1974), and Fama (1998), which forms a date-by-date portfolio of all event firms and regresses it on factor returns. Section 2 writes down each of these classical estimands in our notation and identifies the parameter restriction under which our specification reproduces it.

In parallel, the asset-pricing literature has developed alternative ways to construct the counterfactual return that an event firm would earn in the absence of the event. Daniel et al. (1997) match event firms to characteristic-bucketed nonevent benchmarks, a discrete-bucket approach that anticipates the conditional latent factor models that emerged later, including the instrumented PCA of Kelly et al. (2019), the large-panel conditional risk-premium estimator of Gagliardini et al. (2016), the regressed-PCA method of CRW, on which our Stages 1–3 build, the autoencoder asset-pricing model of Gu et al. (2021), the arbitrage portfolios of Kim et al. (2020), and the projected PCA work of Fan et al. (2016) and Fan et al. (2022). The common feature is that the conditional alpha and the conditional factor loadings depend parametrically or semiparametrically on firm characteristics. Our baseline component sits in this tradition; the firm-event and macro-event components are the new objects that the framework adds. On the inference side, Giglio and Xiu (2021) recover the risk premium of a candidate factor from its projection onto the span of latent factors, and Giglio et al.

(2021) develop alpha tests against estimated latent-factor benchmarks; the loading-span projection of the macro premium and the event-alpha inference below are their event-calendar counterparts.

A third strand focuses on returns around market-wide announcements. The announcement-day premium documented by Savor and Wilson (2014) is reinterpreted in our model as the unconditional aggregate of a characteristic-driven macro-event premium; Lucca and Moench (2015) document the pre-FOMC drift, and Ai and Bansal (2018) connect the announcement premium to consumption-based asset pricing. The within-firm recurring-event differential of Hartzmark and Solomon (2018) similarly corresponds to a cross-sectional average of the total firm-event premium. The framework identifies heterogeneity in that aggregate.

A related literature studies how firm-level news travels across firms. Earnings releases can convey information about industry peers (Foster, 1981); customer–supplier links, shared analyst coverage, and other economic connections can generate predictable cross-firm return patterns (Cohen and Frazzini, 2008; Ali and Hirshleifer, 2020). More generally, such spillovers require the empirical design to specify which firms are exposed to which events (Berg et al., 2021). Our application uses a pre-specified industry-peer mapping and asks whether conditional returns differ on the trading day after a large peer's earnings announcement.

Recent econometric work has refined the inferential foundations of event-study designs in panel data with discrete treatment timing. Borusyak et al. (2024), de Chaisemartin and D'Haultfœuille (2020), and Roth et al. (2023) develop robust estimators for staggered difference-in-differences panels under heterogeneous treatment effects, and Chernozhukov et al. (2018) provides tools for estimating heterogeneous treatment effects with high-dimensional conditioning variables. Closest to our motivation, Goldsmith-Pinkham and Lyu (2025) show that when the benchmark factor model is misspecified, classical event-study estimators are inconsistent for the causal effect of the event—most severely at long horizons, in volatile periods, and when event timing comoves with market conditions—and propose synthetic-control replicating portfolios that dispense with a factor structure altogether. We address the same concern within a model-based design: the conditional benchmark is not fixed ex ante but estimated date by date from non-event firms, jointly with the event premia, although its characteristic-based factor structure remains a maintained specification. Our scheduling condition is only a measurability restriction and is neither stronger nor weaker than unconfoundedness. The separate potential-outcome conditional-mean restrictions deliver model-based identification of the heterogeneous firm-event and macro-event coefficients. Generated-regressor uncertainty then requires the variance corrections developed below.

The paper proceeds as follows. Section 2 sets up the potential outcomes, introduces the data generating process, and develops the connections to the classical event-study estimands. For each of the six—cumulative abnormal returns, multivariate regression model coefficients, calendar-time portfolio alphas, characteristic-matched benchmarks, macro-announcement premia, and recurring firm-event premia—we write down the population analogue, identify the parameter restriction in our specification that recovers it, and discuss what the specification and the estimator each add.

Section 2.4 states the multiple-event and peer-exposure version of the model that the empirical work uses. Section 3 establishes identification of the parameters, presents the four-stage estimator, and summarizes the large-sample theory and the bootstrap we use for standard errors. Section 4 applies the estimator to the CRSP universe of U.S. common stocks over 1980–2024. Four Internet Appendices follow. Appendix A reports the Monte Carlo simulation. Appendix B states the joint limit theorem for the four-stage estimator and the bootstrap result. Appendix C collects the proofs. Appendix D lists the firm characteristics.

2 Setting

2.1 Potential outcomes

Let R_{it} be firm i's return on date t, $Z_{i,t-1} \in \mathbb{R}^L$ its predetermined characteristics, and $D_{i,t-1}$ and $D_{m,t-1}$ indicators for a scheduled firm event and a scheduled macro announcement. The four potential returns are $R_{it}(d, m)$, $d, m \in \{0, 1\}$, and the observed return satisfies

$$R_{it} = R_{it}(D_{i,t-1}, D_{m,t-1}). \tag{1}$$

Throughout, $\mathcal{I}_{t-1}$ denotes the information set generated by returns, characteristics, and event indicators observed through date $t-1$; conditional moments in this section are taken with respect to it. The event calendar is pre-scheduled, so both indicators are known by $t-1$ and are $\mathcal{I}_{t-1}$-measurable.

One firm event and one macro event carry the exposition. That is a readability choice, not a restriction of the framework: firms have several kinds of scheduled news, several macro series are released on overlapping calendars, and a firm's return can respond to its industry peers' announcements as well as to its own. Section 2.4 accommodates all three, replacing the two binary indicators by J firm-event types, M macro types, and an industry peer-announcement indicator, and it is that version the limit theory of Section 3.3, the simulations of Internet Appendix A, and the estimates of Section 4 all use. The model below is its $J = M = 1$ case with the peer channel switched off. Keeping the subscripts implicit until then makes the identification argument readable; the general version changes the dimension of the event vector, not the argument.

Writing the potential return in (1) as a function of the date-$(t-1)$ indicators alone is itself a restriction, and it is worth naming what it assumes (Rubin, 1974).

(i) *No anticipation.* The date-t return responds to an event only from the date the event occurs. Pre-scheduling is consistent with this restriction: the date of a release is public well in advance, its content is not. What the restriction rules out is a drift before the event that is attributable to the event itself, which would contaminate the non-event firm-dates that serve as controls, and carryover from one announcement into a later one-day window. A response

that unfolds over several days is accommodated by dating the indicator, as with the next-day peer-announcement indicator of Section 2.4, rather than by widening the return window.

(ii) *No interference beyond the included indicators.* A firm's return does not depend on which other firms announce, except through the event indicators that the specification already contains (Rubin, 1980). In the baseline model of Section 2.2 those indicators are the firm's own event status and the macro calendar, so other firms' announcements do not enter. Section 2.4 appends an indicator for whether sufficiently many industry peers announce, and Section 4.2.3 estimates its coefficient. Peer effects of any magnitude are then permitted; the restriction is only that they operate through that indicator (Berg et al., 2021).

Pre-scheduling and these restrictions do not by themselves deliver exogeneity. The conditional-mean restrictions of Section 3.1, rather than the calendar, supply the identifying content.

The firm-event effect on a non-announcement date and the macro effect for a non-event firm are

$$\tau^{\gamma}(z) := \mathbb{E}[R_{it}(1,0) - R_{it}(0,0) \mid Z_{i,t-1} = z], \tag{2}$$

$$\tau^{\delta}(z) := \mathbb{E}[R_{it}(0,1) - R_{it}(0,0) \mid Z_{i,t-1} = z], \tag{3}$$

$$\tau^{\gamma\delta}(z) := \mathbb{E}[R_{it}(1,1) - R_{it}(0,1) \mid Z_{i,t-1} = z]. \tag{4}$$

The first two are clean comparisons, each switching one event on while the other is off. The third is the firm-event effect *on* an announcement date: it completes the two-way design and is carried by the coincidence cell, whose level and exposure coefficients appear in the return model below. The comparison of $\tau^{\gamma\delta}$ with τ^{γ} is where the economics of competing news lives: if macro announcements crowd out attention to firm-specific news, firm-event compensation should be lower on announcement days (Hirshleifer et al., 2009); if firm news arriving together with macro news compounds announcement risk, it should be higher (Savor and Wilson, 2014; Ai and Bansal, 2018). Returns on coincidence dates, dates on which a firm's own event falls on a macro announcement, are modeled flexibly so that neither clean comparison is forced to absorb the other event.

All three are functions of a characteristic value z, held fixed, rather than of a particular firm on a particular date. Firm i's own effect on date t is the same function evaluated at $Z_{i,t-1}$, and it does move over time, because the firm's characteristics move while the model's coefficients do not. Holding z fixed is what isolates the parameters from the composition of the cross-section, and it is why none of the definitions carries a date subscript.

The causal reading of these contrasts is not uniform across event types. The macro releases are verifiably pre-scheduled, the calendar is set months in advance, so τ^{δ}, and the announcement-day conditioning in $\tau^{\gamma\delta}$, carry the cleanest causal reading, subject to the conditional-mean restrictions below. Earnings dates are firm-chosen recorded dates: a firm that shifts bad news moves itself across cells, so the earnings component of τ^{γ} is a conditional-return differential on recorded dates rather than the effect of an exogenously assigned announcement; the on-time/delayed split of Section 4.2

measures that margin and is read as the boundary of the causal interpretation. Distribution ex-dates are mechanical: the ex-day price adjustment the framework recovers is a tax-and-microstructure fact, not an announcement effect, and is labeled accordingly throughout. Restriction (ii) excludes same-day industry information transfer; the next-day peer estimate of Section 4.2.3 shows that channel is live at the adjacent horizon, so the same-day exclusion is a maintained restriction rather than an established fact.

We are interested in the composition of $\tau^{\gamma}(z)$ as much as in its size. An announcement can raise a firm's expected return because it carries compensation unrelated to systematic risk, or because it changes the firm's exposure to the common factors and so changes the risk premium the firm earns. The return model separates the two channels, for each event type and across firms.

2.2 Return model

Returns co-move through K latent factors. A firm event is news about one firm, but a macro release is news about the market, so it can move the common factor itself rather than a single firm's return. The factor therefore needs potential values of its own: let $F_t(0)$ and $F_t(1)$ be the date-t factor without and with a macro announcement, and write

$$F_t(1) = F_t(0) + \lambda + \eta_t, \qquad \mathbb{E}[\eta_t \mid \mathcal{I}_{t-1}] = 0, \qquad F_t = F_t(D_{m,t-1}). \tag{5}$$

An announcement moves the factor by a predictable amount and an unpredictable one. The vector λ is the average factor shift on announcement dates, the *factor announcement premium*, the object to which a factor-based account of the macro premium would attribute that premium. The surprise η_t is unforecastable given date-$(t-1)$ information and so carries no compensation. The last equality in (5) is consistency for the factor, the analogue of (1) for returns: the realized factor is the potential factor evaluated at the realized announcement status. Only one of $F_t(0)$ and $F_t(1)$ is ever observed on a given date.

Given (5), the four return cells are

$$R_{it}(0,0) = Z'_{i,t-1}\alpha + Z'_{i,t-1}BF_t(0) + U_{it}(0), \tag{6}$$
$$R_{it}(1,0) = Z'_{i,t-1}(\alpha+\gamma) + Z'_{i,t-1}(B+\Gamma)F_t(0) + U_{it}(1), \tag{7}$$
$$R_{it}(0,1) = Z'_{i,t-1}\alpha + Z'_{i,t-1}BF_t(1) + U_{it}(0), \tag{8}$$
$$R_{it}(1,1) = Z'_{i,t-1}(\alpha+\gamma_a) + Z'_{i,t-1}(B+\Gamma_a)F_t(1) + U_{it}(1). \tag{9}$$

These are the four combinations of the two indicators, and taking them in order shows what each event contributes. With neither event, (6) is an ordinary conditional factor model in which characteristics carry both the level α and the loadings B. A firm event alone, (7), changes two things: the level, by γ, and the factor exposure, by Γ. Systematic exposures move with firm-specific news arrival (Patton and Verardo, 2012; Savor and Wilson, 2016), which is why Γ enters:

with $\Gamma = 0$ an event could still carry compensation, but none of it could be attributed to a change in the firm's exposure to common factors. A macro announcement alone, (8), leaves the firm's own coefficients untouched and instead replaces $F_t(0)$ by $F_t(1)$, which is the channel (5) describes. When both occur, (9), the coefficients (γ_a, Γ_a) are free to differ from (γ, Γ), so the coincidence date is not forced to equal the sum of the two separate effects; this is what keeps a simultaneous macro release from being mechanically charged to the firm-event effect. Each of the two single event cases corresponds to an established literature, and Section 2.3 maps those literatures onto the parameter restrictions that recover their estimands. The both-at-once case corresponds to none of them, and isolating it is one of the things the joint specification buys.

Equations (6)–(9) are written in $F_t(0)$ and $F_t(1)$, and only one of the two is realized on any given date. A single equation that holds on every date must therefore be written in objects the panel delivers. Define

$$\bar{F}_t := F_t - \lambda D_{m,t-1}, \qquad \delta := B\lambda, \qquad c := \gamma_a + \Gamma_a \lambda. \tag{10}$$

The first is the realized factor with the announcement shift stripped out: $\bar{F}_t = F_t(0)$ on a date without a macro release and $F_t(0) + \eta_t$ on a date with one. It is one series, defined on every date, playing the role that $F_t(0)$ plays in (6). Removing λ from the factor does not remove its effect on returns: on an announcement date that effect is $Z'_{i,t-1} B\lambda$, and δ names it. Where λ was a shift in a latent factor, δ is a shift in expected returns attached to an indicator the calendar fixes in advance, which is what makes it estimable. The third definition applies the same substitution to the coincidence cell, the both-at-once case (9), a firm event on a macro-announcement date, so c is that cell's intercept measured against $\bar{F}_t$.

Substituting (10) into the four cells collapses them into one equation, the observed-data model

$$\begin{aligned} R_{it} = Z'_{i,t-1}\Big[&\alpha + B\bar{F}_t + \delta D_{m,t-1} + (\gamma + \Gamma \bar{F}_t) D_{i,t-1}(1 - D_{m,t-1}) \\ &+ (c + \Gamma_a \bar{F}_t) D_{i,t-1} D_{m,t-1}\Big] + U_{it}. \end{aligned} \tag{11}$$

Every date now contributes to the same regression. The factor enters as the single series $\bar{F}_t$, and the remaining parameters are attached to indicators that are fixed before the return is realized. The realized innovation is

$$U_{it} := U_{it}(D_{i,t-1}), \qquad \mathbb{E}[U_{it}(d) \mid \mathcal{I}_{t-1}, \bar{F}_t] = 0, \quad d \in \{0, 1\}. \tag{12}$$

Thus macro news enters through the common factor under (5), while the idiosyncratic innovation may differ between firm-event and ordinary dates.

We impose one further restriction on the factor: its recentered conditional mean is constant, $\mathbb{E}[\bar{F}_t \mid \mathcal{I}_{t-1}] = \mu$, and by iterated expectations $\mu = \mathbb{E}[\bar{F}_t]$. The model therefore carries no predictor of the factor mean. Identifying and estimating time variation in a factor's mean is a separate

problem (Gagliardini et al., 2016), orthogonal to the decomposition of event returns studied here. The identifying content of this restriction for δ is only that the counterfactual factor mean not differ across announcement status: announcements are not scheduled on information about the factor level. The conditional-on-history form is the strengthening that the time-series estimation theory uses, it makes the calendar score a martingale difference (Internet Appendix B.2), and we impose it once for both purposes. Under this restriction the model implies

$$\tau^{\gamma}(z) = z'(\gamma + \Gamma\mu), \qquad \tau^{\delta}(z) = z'\delta, \qquad \tau^{\gamma\delta}(z) = z'(c + \Gamma_a\mu). \tag{13}$$

The first is the decomposition the paper is after: $z'\gamma$ is the *event alpha*, the part of the expected event return that is not compensation for factor exposure, and $z'\Gamma\mu$ is the *event risk premium*, the part earned because the announcement changes the firm's loadings. Integrating out the factor realization replaces $\bar{F}_t$ by μ; on an announcement date $F_t(0)$ and $\bar{F}_t$ differ by η_t, whose conditional mean is zero by (5).

The macro effect admits no analogous split: under (10) it is factor compensation in its entirety, $z'\delta = z'B\lambda$, the return counterpart of the mean factor shift λ.

The third contrast splits like the first: $\tau^{\gamma\delta}(z) = z'\{\gamma_a + \Gamma_a(\mu + \lambda)\}$ is an announcement-day event alpha $z'\gamma_a$ plus compensation for the announcement-day exposure Γ_a, evaluated at the factor mean the announcement induces. Its reduced form $z'(c + \Gamma_a\mu)$ is identified without further restrictions; separating γ_a from the compensation uses the same maintained factor channel that delivers λ.

2.3 Existing event-study estimands as special cases

The model nests the principal estimands used in event studies. This connection clarifies what is gained by allowing event returns to vary with firm characteristics, latent factor realizations, and the broader announcement calendar.

The cumulative abnormal return (CAR) of Fama et al. (1969), Brown and Warner (1985), and MacKinlay (1997) fits a market model $R_{it} = \alpha_i + \beta_i R_{M,t} + \epsilon_{it}$, with $R_{M,t}$ the observed market return, and treats the residual over an event window as the abnormal return. Our model reduces to this benchmark when the single factor is observed and equal to $R_{M,t}$, the event effect is constant across firms, and the announcement channel is shut ($\Gamma = 0$, $\lambda = 0$ and hence $\delta = 0$, with the coincidence-cell coefficients equal to their ordinary-date counterparts, $\gamma_a = \gamma$ and $\Gamma_a = \Gamma = 0$). It generalizes the benchmark by estimating the factor from the full panel and allowing both factor exposures and event premia to vary with firm characteristics. The term $\Gamma\bar{F}_t$ further allows the event premium to depend on the realized factors. Goldsmith-Pinkham and Lyu (2025) show that misspecification of the benchmark model makes CAR-based estimates of the event effect inconsistent; estimating the benchmark within each date from the non-event cross-section, rather than from each firm's own pre-event time series, is this paper's response to that critique.

The multivariate regression (MVR) model of Schipper and Thompson (1983) and Binder (1985)

pools event firms in a single panel regression $R_{it} = \alpha_i + \beta_i' f_t + g\, D_{i,t-1} + \epsilon_{it}$, with observed factors f_t and a common event coefficient g. It corresponds to the restrictions that the factor is observed, the event effect is constant across firms, and $\Gamma = 0$, with the announcement channel shut ($\lambda = 0$, hence $\delta = 0$, and $\eta_t \equiv 0$) and the two firm-event coefficient maps equal across announcement statuses ($\gamma_a = \gamma$, $\Gamma_a = \Gamma$). We replace firm fixed effects with a parsimonious characteristic representation and estimate heterogeneous event premia jointly with the latent factor structure.

The calendar-time portfolio (CT) approach of Jaffe (1974), Mandelker (1974), and Fama (1998) forms a portfolio of all firms with an event on each date and regresses that portfolio return on factor returns. Its intercept combines the baseline returns of event firms, the firm-event premium, and any macro-announcement premium earned on the same dates. Our model separates these components and recovers their cross-sectional heterogeneity rather than a single portfolio average.

The characteristic-matched benchmark (DGTW) approach of Daniel et al. (1997) matches each event firm to a portfolio of nonevent firms with similar size, book-to-market, and momentum, and defines the abnormal return as the firm–benchmark difference $\mathrm{AR}_{it}^{\mathrm{DGTW}} = R_{it} - R_{\mathcal{B}(Z_{i,t-1}),t}$. The benchmark portfolio return estimates the conditional mean return given the characteristic bucket, $Z_{i,t-1}'(\alpha + B\bar{F}_t)$ on macro-free dates ($Z_{i,t-1}'\delta\, D_{m,t-1}$ is added on announcement dates), under a step-function specification of the loadings. Our framework replaces discrete buckets with a continuous characteristic specification and separates baseline returns, factor exposures, firm-event effects, and macro-event effects.

The macro-announcement-day comparison (SW) of Savor and Wilson (2014) contrasts the unconditional mean return on FOMC, employment, and inflation announcement days with non-event days,

$$\Delta_m \;:=\; \mathbb{E}[R_{M,t} \mid D_{m,t-1} = 1] - \mathbb{E}[R_{M,t} \mid D_{m,t-1} = 0],$$

with $R_{M,t}$ the observed market return as above. This contrast is an aggregate version of $z'\delta$. Our framework estimates the full vector δ, revealing which firm characteristics account for the announcement-day premium. Related work studies the pre-FOMC drift and consumption-based explanations (Lucca and Moench, 2015; Ai and Bansal, 2018).

The recurring-firm-event comparison (HS) of Hartzmark and Solomon (2018) documents predictable abnormal returns on recurring firm-event dates by comparing event-day returns to non-event-day returns within firm, $\Delta^{\mathrm{HS}}(z) = \mathbb{E}[R_{it} \mid D_{i,t-1} = 1, Z_{i,t-1} = z] - \mathbb{E}[R_{it} \mid D_{i,t-1} = 0, Z_{i,t-1} = z]$. When firm and macro events are independent, the factor is mean-independent of firm-event timing, and coincidence cells carry negligible weight, the within-firm differential equals $z'\gamma + z'\Gamma\mu$ for a firm with characteristics z. We decompose this average effect into an event level and a factor-dependent component, linking recurring-event predictability to the event-risk mechanism studied by Savor and Wilson (2016).

2.4 Multiple event types and peer exposure

The empirical specification estimates J firm-event or exposure types and M macro-announcement types jointly. Write $D^{(j)}_{i,t-1}$ for the indicator of a type-j firm event and $D^{(g)}_{m,t-1}$ for the indicator of a type-g macro release, both predetermined at $t-1$, and collect them in

$$\boldsymbol{D}^{f}_{i,t-1} := (D^{(1)}_{i,t-1}, \ldots, D^{(J)}_{i,t-1})', \qquad \boldsymbol{D}^{m}_{t-1} := (D^{(1)}_{m,t-1}, \ldots, D^{(M)}_{m,t-1})'.$$

The scalars of Section 2.1 are the case $J = M = 1$. Let

$$D_{m,t-1} := \max_{g \le M} D^{(g)}_{m,t-1}$$

record whether any macro release occurs. The potential factor response is additive across macro types: writing $\mathbf{m} \in \{0,1\}^M$ for a value of $\boldsymbol{D}^{m}_{t-1}$, $F_t(\mathbf{m}) = F_t(\mathbf{0}) + \sum_{g=1}^{M} m_g(\lambda_g + \eta_{g,t})$ with $\mathbb{E}[\eta_{g,t} \mid \mathcal{I}_{t-1}] = 0$ and $F_t = F_t(\boldsymbol{D}^{m}_{t-1})$, so that $\bar{F}_t = F_t(\mathbf{0}) + \sum_{g=1}^{M} \eta_{g,t} D^{(g)}_{m,t-1}$, and the structural factor-response model implies $\boldsymbol{\delta} = B\boldsymbol{\lambda}$, where bold symbols collect type-indexed objects columnwise: $\boldsymbol{\delta} := (\delta_1, \ldots, \delta_M) \in \mathbb{R}^{L \times M}$, $\boldsymbol{\lambda} := (\lambda_1, \ldots, \lambda_M) \in \mathbb{R}^{K \times M}$, hats inherited entrywise, a convention used throughout, bold marking collection across event types rather than vector dimension B, common to all types, is never bold. These are the multivariate counterparts of the first two definitions in (10), and they fix the factor response on dates carrying more than one release. The number of macro types is unrestricted relative to K, and with $M > K$ the model implies $\operatorname{rank}(\boldsymbol{\delta}) \le K$: the M premium vectors are linearly dependent, an implication of the maintained restriction $\boldsymbol{\delta} = B\boldsymbol{\lambda}$. The single-event model extends by stacking coefficient blocks:

$$\begin{aligned} R_{it} = Z'_{i,t-1}\Big[\alpha + B\bar{F}_t + \sum_{g=1}^{M} \delta_g D^{(g)}_{m,t-1} + \sum_{j=1}^{J} (\gamma_j + \Gamma_j \bar{F}_t) D^{(j)}_{i,t-1}(1 - D_{m,t-1}) \\ + \sum_{j=1}^{J} \Big(\sum_{g=1}^{M} c_{jg} D^{(g)}_{m,t-1} + \Gamma_{a,j} \bar{F}_t \, D_{m,t-1} \Big) D^{(j)}_{i,t-1} \Big] + U_{it}. \end{aligned} \tag{14}$$

Every date now enters (14) exactly once. A firm-event observation is described by the clean coefficients (γ_j, Γ_j) when no release occurs and by the coincidence coefficients $(c_{jg}, \Gamma_{a,j})$ when at least one does. As in the single-event case, c_{jg} is an increment over δ_g, not a total: a type-j announcing firm on a date carrying only release g has conditional mean $Z'_{i,t-1}(\alpha + B\bar{F}_t + \delta_g + c_{jg} + \Gamma_{a,j}\bar{F}_t)$, so c_{jg} measures it against a non-announcing firm on the same date and $c_{jg} = \gamma_{a,jg} + \Gamma_{a,j}\lambda_g$: the multivariate counterpart of the last definition in (10), where $\gamma_{a,jg}$ is the structural announcement-day level of the cell. The innovation is $U_{it} := U_{it}(\boldsymbol{D}^{f}_{i,t-1})$ with $\mathbb{E}[U_{it}(\mathbf{d}) \mid \mathcal{I}_{t-1}, \bar{F}_t] = 0$ for every $\mathbf{d} \in \{0,1\}^J$, the values taken by $\boldsymbol{D}^{f}_{i,t-1}$.

The specification is saturated on the configurations the empirical work uses. The clean coefficients (γ_j, Γ_j) are free for every type j on release-free dates, and the coincidence level c_{jg} is free for every pair (j, g) on dates that carry exactly one release. Four pooling restrictions govern the

remainder, since a free coefficient pair for each release configuration is not estimable from the few dates on which two releases land together. Levels add across macro types; $\Gamma_{a,j}$ does not vary with g; the announcement-day exposure loads once whenever $D_{m,t-1} = 1$ rather than accumulating with the number of releases; and effects add across firm-event types, so a firm with two kinds of scheduled news on one day earns $\gamma_j + \gamma_{j'}$, an interaction being available only by adding $D_{i,t-1}^{(j \cap j')}$ as a further indicator with its own coefficient pair. The third and first interact: because c_{jg} bundles a per-date object with a per-release one, summing it over active releases charges the announcement-day event intercept once per release. Note finally that $1 - D_{m,t-1} = \prod_g (1 - D_{m,t-1}^{(g)})$, so the maximum in the clean-event switch is notation for "no release of any type" and imposes nothing; the content of $D_{m,t-1}$ sits in the exposure term.

The multiple-event identification and limit results follow by stacking a fixed number of design blocks, subject to the corresponding joint-rank condition. No new economic theorem is required.

Peer exposure is one additional firm-level indicator. Let $\text{ind}(i)$ denote the set of firms sharing firm i's industry classification, $D_{k,t-1}^{(\bullet)}$ the indicator of the event type whose spillover is at issue, and $k_0 \geq 1$ a threshold count. An industry-peer exposure indicator is

$$D_{i,t-1}^{(\text{spill})} := \mathbf{1}\left\{ \sum_{k \in \text{ind}(i) \setminus \{i\}} D_{k,t-1}^{(\bullet)} \geq k_0 \right\}, \tag{15}$$

which marks the dates on which at least k_0 of firm i's peers carry that event. Firm i is excluded from the sum, so the indicator is not a relabelling of its own event status, and it is predetermined at $t-1$ whenever its peers' calendars are. It therefore enters (14) as one more component of $\boldsymbol{D}_{i,t-1}^{f}$, with its own level and exposure blocks, and its coefficient is the conditional-return difference associated with that exposure mapping. A causal spillover interpretation additionally requires that the mapping and controls exclude common shocks. The peer set can instead be built from supply-chain links (Cohen and Frazzini, 2008), characteristic similarity, or shared analyst coverage (Ali and Hirshleifer, 2020); the choice is application-specific. Section 4.2.3 takes $\text{ind}(i)$ to be the three-digit SIC industry, restricts the source firms to bellwethers, sets $k_0 = 1$ with earnings as the source event, and dates the indicator one trading day after the peer announcement, which measures delayed rather than contemporaneous exposure.

For readability, the remaining derivations use $J = M = 1$. The empirical estimator stacks all event types, block by block, as Section 3.2 sets out.

3 Identification, Estimation, and Inference

This section establishes identification and presents the four-stage estimator. The estimation stages are the sample counterparts of the conditional moments that carry the identification argument. The closing subsection summarizes the limit theory and the bootstrap behind every reported standard error.

3.1 Identification

Recall $\mathcal{I}_{t-1}$ from Section 2.1, and let $\mathcal{G}_{t-1} := \sigma(\mathcal{I}_{t-1}, \bar{F}_1, \ldots, \bar{F}_{t-1})$ add the history of the latent factor. Two sets of conditions carry the identifying content; their formal statements are Assumptions 1 and 2 in Internet Appendix B.2.

The first set governs sampling and timing. Event indicators and characteristics are predetermined, and the innovations and the recentered factor satisfy the conditional-mean restrictions

$$\mathbb{E}[U_{it}(d) \mid \mathcal{G}_{t-1}, \bar{F}_t] = 0, \qquad \mathbb{E}[\bar{F}_t \mid \mathcal{G}_{t-1}] = \mu, \quad d \in \{0, 1\}.$$

Both displays carry economic content. The first makes the characteristic-based benchmark complete: common return variation enters through $Z'_{i,t-1}B\bar{F}_t$, and what remains is conditionally mean zero given the factor realization and the date-$(t-1)$ information. The second extends the constant-factor-mean restriction of Section 2.2 from $\mathcal{I}_{t-1}$ to $\mathcal{G}_{t-1}$: the factor's own history, like the announcement calendar, does not forecast the recentered factor mean. The remaining conditions of Assumption 1 are regularity: on every date some firms carry no firm event, these are the *controls*, and their realized characteristics span, so the control design is nonsingular; each release calendar is nonconstant, announcement and clean dates both occur, and the clean-event design has full rank. The control group is not static: membership is determined date by date, so the same firm serves as a control on dates without its own event and as an event firm on its event dates. With several event types the same conditions are read for the stacked objects.

The second set, Assumption 2, governs the factor structure and contains the identifying normalization. The number of factors K is fixed and known, the loading matrix has rank K, and the centered factor second moment $\Sigma_F := \lim_{T\to\infty} T^{-1} \sum_{t\le T} \mathbb{E}[(\bar{F}_t - \mu)(\bar{F}_t - \mu)']$ exists and is positive definite. These conditions identify the loading space. The key identifying assumption, which separates alpha from risk premia, is the orthogonality normalization

$$\alpha' B = 0$$

(Chen et al., 2025). The panel identifies the sum of the baseline level and average factor compensation, not the two separately: a constant can be moved from the factor mean into the intercept without changing any conditional return. Orthogonality resolves that ambiguity, at the baseline and at events alike.

Taking K as known deserves comment. Once returns are projected into the L characteristic-managed portfolios, standard factor-model intuition applies to the slope path, and eigenvalue-based criteria (Bai and Ng, 2002; Onatski, 2010; Ahn and Horenstein, 2013) could in principle select K. In our sample they do not settle it: the spectrum has one dominant factor and no elbow thereafter, and the eigenvalue-ratio criterion selects $K = 1$ (Section 4.2). We therefore treat $K = 3$ as a working choice, on which all reported standard errors condition, report the headline premia at $K \in \{2, 4, 5\}$, and measure the cost of over-extraction in Internet Appendix A.

Identification proceeds in two steps, stated in conditional population moments: on each date the conditional moments of the control cross-section pin down the characteristic slope $a_t := \alpha + B\bar{F}_t + \delta D_{m,t-1}$, and the conditional means and centered second moments of the slope path then separate the parameters; the moments may vary freely over time.

What the panel identifies without any normalization is the reduced form: the macro premia, the loading space through its projector $P_B := B(B'B)^{-1}B'$, the average slope $\bar{a} = \alpha + B\mu$, the average firm-event effect $\gamma + \Gamma\mu$ and its announcement-day counterpart $c + \Gamma_a\mu$, the centered factor components $B(\bar{F}_t - \mu)$ and $\Gamma(\bar{F}_t - \mu)$, and every fitted conditional return, including, on a full-rank coincidence design, those on the coincidence cell, but not α, γ, and μ separately. It cannot do more because, for any $\zeta \in \mathbb{R}^K$, the translation

$$\bar{F}_t^\star = \bar{F}_t - \zeta, \qquad \alpha^\star = \alpha + B\zeta, \qquad \gamma^\star = \gamma + \Gamma\zeta$$

reproduces every conditional mean in the model. The restriction $\alpha'B = 0$ selects one member of this class, the unique translation achieving it is $\zeta = -(B'B)^{-1}B'\alpha$, and delivers the split: because $P_B\alpha = 0$, the conditional alpha $\alpha = (I_L - P_B)\bar{a}$ and the average factor risk premium $B\mu = P_B\bar{a}$ are orthogonal components of the single identified object $\bar{a}$, and γ, $\Gamma\mu$, $B\bar{F}_t$, and $\Gamma\bar{F}_t$ follow, none depending on the factor coordinate system. The coordinate conventions add the labels: B, Γ, and $\{\bar{F}_t\}$ become identified individually, up to a common sign in each factor column. On the same terms, a full-rank coincidence design identifies (c, Γ_a), and under the factor-response model the factor path, the structural coincidence intercept $\gamma_a = c - \Gamma_a\lambda$, and the factor announcement premium $\lambda = B'\delta$ are identified as well. One boundary remains under every normalization: on an announcement date the counterfactual factor $F_t(0)$ and the announcement surprise η_t of (5) cannot be separated. Proposition 1 in Internet Appendix B.2 states all of this formally, and Internet Appendix C proves it.

The division is relative to the factor model the panel admits. What the split calls the conditional alpha is the part of the average characteristic slope that the K extracted factors do not price, so a source of compensated risk whose loadings lie in the characteristic span but which is not among those K is carried by α rather than by $B\mu$. Alpha here is an abnormal return measured against a stated benchmark, not a model-free measure of mispricing (Giglio et al., 2021).

The coordinate conventions are almost free. For any nonsingular $K \times K$ matrix Q, the representation

$$(BQ^{-1},\ Q\bar{F}_t,\ Q\mu,\ \Gamma Q^{-1})$$

leaves P_B, $B\bar{F}_t$, $B\mu$, γ, $\Gamma\bar{F}_t$ and $\Gamma\mu$ unchanged, so $B'B = I_K$ and an ordered diagonal Σ_F merely pick the PCA representative inside the identified loading space, and what they buy is the ability to label individual factor columns. The one part that is not free is that the diagonal entries be distinct: no rotation delivers a strictly decreasing diagonal when two factors carry equal centered variance, and it is distinctness that identifies the columns of B rather than only $\mathrm{col}(B)$.

One qualification. The condition $\alpha' B = 0$ makes two vectors of characteristic coefficients orthogonal, so it depends on the scale on which the characteristics are measured: rescaling one characteristic relative to the others would shift part of the average slope between the conditional alpha and the risk premium. We rank-normalize every characteristic to a common scale, which fixes that choice. The same argument applies to the basis and not only the scale: α and $B\mu$ are orthogonal components of a slope vector on one fixed set of characteristics, so adding or dropping a characteristic moves mass between them.

3.2 The Four-Stage Estimator

We estimate the model's parameters, $(\lambda, \alpha, B, \{\bar{F}_t\}, \gamma, \Gamma, \gamma_a, \Gamma_a)$, in four stages, that replace the identifying conditional population moments with their sample analogs; the fourth comprises companion regressions on the two event cells, Stage 4a for the clean pair (γ, Γ) and Stage 4b for the coincidence pair (c, Γ_a). The stages produce the reduced-form coefficients of (11); the structural parameters follow exactly, λ by minimum distance from $\delta = B\lambda$ and γ_a from $c = \gamma_a + \Gamma_a\lambda$, both of (10).

First, the index sets the stages run on. A control firm is one with no firm event of any *included* type, so with J firm-event types write $D_{i,t-1} := \max_{j \le J} D^{(j)}_{i,t-1}$ and

$$\mathcal{T}^{(0)}_t := \{i : \ D_{i,t-1} = 0\},$$

with $Z^{(0)}_{t-1}$ and $R^{(0)}_t$ stacking characteristics and returns over $\mathcal{T}^{(0)}_t$. Taking the maximum over types is what keeps the control cross-section free of event firms. Pooling over dates, the event observations split by whether a macro release falls on the same date, because (11) assigns the two cases different coefficients:

$$\mathcal{T}^{(1)} := \{(i,t) : D_{i,t-1} = 1\}, \quad \mathcal{T}^{(1,0)} := \{(i,t) \in \mathcal{T}^{(1)} : D_{m,t-1} = 0\}, \quad \mathcal{T}^{(1,m)} := \mathcal{T}^{(1)} \setminus \mathcal{T}^{(1,0)}. \tag{16}$$

We call $\mathcal{T}^{(1,0)}$ the clean-event index set and write $n_1 := |\mathcal{T}^{(1,0)}|$ for its size, n_{1t} for the number of clean-event observations on date t, so that $n_1 = \sum_{t=1}^{T} n_{1t}$. For control firms the model of (11) gives

$$\mathbb{E}[R_{it} \mid \mathcal{I}_{t-1}, \bar{F}_t, D_{i,t-1} = 0] = Z'_{i,t-1} a_t, \qquad a_t := \alpha + B\bar{F}_t + \delta D_{m,t-1}, \tag{17}$$

and a_t is what Stage 1 estimates.

Stage 1 (slope on control firms). For each t, fit the cross-sectional OLS of $R^{(0)}_t$ on $Z^{(0)}_{t-1}$:

$$\hat{a}_t \ := \ (Z^{(0)\prime}_{t-1} Z^{(0)}_{t-1})^{-1} Z^{(0)\prime}_{t-1} R^{(0)}_t, \qquad t = 1, \ldots, T. \tag{18}$$

Stage 1 is run date by date and is the only stage that uses firms with no event of any included type; Stages 2 and 3 see the data only through the resulting slope path, and Stages 4a and 4b use

event observations alone.

Stage 2 (macro-event effect δ).

$$\hat{\delta} \;=\; \frac{\sum_{t=1}^{T}(D_{m,t-1}-\bar{D}_m)\,(\hat{a}_t-\bar{\hat{a}})}{\sum_{t=1}^{T}(D_{m,t-1}-\bar{D}_m)^2}, \qquad \hat{a}_t^0 \;:=\; \hat{a}_t - \hat{\delta}\,D_{m,t-1}. \tag{19}$$

Stage 2 compares the characteristic slopes on macro-announcement and non-announcement dates. Although the factor is still latent at this stage, Assumption 1 implies that the event calendar does not forecast its innovation. Factor realizations therefore affect the precision of $\hat{\delta}$, but not its consistency. Stage 2 estimates δ as a free L-vector; the factor-response implication $\delta = B\lambda$ of (10) is imposed only in Stage 3's minimum-distance recovery of λ.

Equation (19) is the case $M = 1$. The empirical specification estimates the M macro premia jointly, and it is this multivariate version that produces every macro number reported below. Stack the Stage-1 slopes and the release calendars as

$$\hat{A} := (\hat{a}_1,\ldots,\hat{a}_T) \in \mathbb{R}^{L\times T}, \qquad \mathbb{D}^m := (\boldsymbol{D}_0^m,\ldots,\boldsymbol{D}_{T-1}^m)' \in \mathbb{R}^{T\times M},$$

so row t of $\mathbb{D}^m$ holds the release indicators predetermined at t. Multivariate Stage 2 is the time-series OLS

$$\widehat{\boldsymbol{\delta}} \;:=\; \hat{A}\,M_T\,\mathbb{D}^m\big(\mathbb{D}^{m\prime}M_T\,\mathbb{D}^m\big)^{-1} \;\in\; \mathbb{R}^{L\times M}, \qquad \hat{a}_t^0 \;:=\; \hat{a}_t - \widehat{\boldsymbol{\delta}}\,\boldsymbol{D}_{t-1}^m, \tag{20}$$

where M_T is the demeaning projection of Section 3.1 and column g of $\widehat{\boldsymbol{\delta}}$ is $\hat{\delta}_g$, the premium of macro type g. The only new requirement is that $\mathbb{D}^{m\prime}M_T\mathbb{D}^m$ be nonsingular: no release calendar is constant or a linear combination of the others and the constant. Because the calendars are not mutually exclusive, some dates carry two releases, as reported in Section 4.1, this Gram matrix is not diagonal, and each $\hat{\delta}_g$ is a partial effect that holds the other release calendars fixed. The later stages use the residualized slopes $\hat{a}_t^0$ of (20), from which all M macro effects have been removed simultaneously. Whatever the value of M, the Stage-2 normal equations make the residualized slopes exactly orthogonal to the demeaned calendars; this exact annihilation is why Stage-2 estimation error does not propagate into the Stage-3 eigenvectors (Internet Appendix C).

Stage 3 (factor structure, alpha, and factor announcement premium λ). Apply PCA to the column-demeaned slopes. Let $\hat{A}^0 := (\hat{a}_1^0,\ldots,\hat{a}_T^0) \in \mathbb{R}^{L\times T}$ and define the eigenvalue decomposition $(\hat{A}^0\,M_T\,\hat{A}^{0\prime}/T)\,\hat{B} = \hat{B}\,V$, with V diagonal. The columns of $\hat{B}$ are the eigenvectors corresponding to the K largest eigenvalues of $\hat{A}^0\,M_T\hat{A}^{0\prime}/T$, normalized so $\hat{B}'\hat{B} = I_K$. Then

$$\hat{F}_t \;=\; \hat{B}'\,\hat{a}_t^0, \qquad \hat{\alpha} \;=\; (I_L - \hat{B}\hat{B}')\,\bar{\hat{a}}^0, \qquad \bar{\hat{a}}^0 := \frac{1}{T}\sum_{t=1}^{T}\hat{a}_t^0, \qquad \hat{\mu} \;:=\; \hat{B}'\,\bar{\hat{a}}^0. \tag{21}$$

Because $\alpha' B = 0$, the projection $\hat{B}'\hat{a}_t^0$ recovers the factor including its mean, in the coordinates Stage 3 delivers, and $\hat{\mu} = T^{-1}\sum_t \hat{F}_t$. Neither $\hat{F}_t$ nor $\hat{\mu}$ is separately comparable to $\bar{F}_t$ and μ; the products $\hat{\Gamma}\hat{F}_t$ and $\hat{\Gamma}\hat{\mu}$ are. Stage 3 also delivers the loading-span projection $\hat{B}'\hat{\delta}$ and the orthogonal component $(I_L - \hat{B}\hat{B}')\hat{\delta}$, columnwise in $\widehat{\boldsymbol{\delta}}$ when Stage 2 is multivariate. The latter is the orthogonal component reported in the descriptive span decomposition of Section 4.2; the former returns as the minimum-distance estimate of the factor announcement premium below.

Because the factor-response model implies $\boldsymbol{\delta} = B\boldsymbol{\lambda}$ (Section 2.4), the Stage-2 and Stage-3 outputs already contain the structural estimate of the factor announcement premium, and no further regression is required: because $\hat{B}'\hat{B} = I_K$,

$$\hat{\boldsymbol{\lambda}} \;:=\; \hat{B}'\widehat{\boldsymbol{\delta}} \;=\; \underset{\boldsymbol{\lambda}\in\mathbb{R}^{K\times M}}{\arg\min} \left\|\widehat{\boldsymbol{\delta}} - \hat{B}\boldsymbol{\lambda}\right\|_F^2$$

is the minimum-distance estimator of $\boldsymbol{\lambda}$; the problem decouples column by column into $\hat{\lambda}_g = \hat{B}'\hat{\delta}_g$. Under the restriction the entire $\sqrt{T}$ sampling variation of $\widehat{\boldsymbol{\delta}}$ lies in $\mathrm{col}(B)$, so $\hat{B}\hat{\boldsymbol{\lambda}}$ and $\widehat{\boldsymbol{\delta}}$ are $\sqrt{T}$-equivalent and the minimum-distance step costs nothing at first order. Entries of $\hat{\boldsymbol{\lambda}}$ are expressed in the Stage-3 factor coordinates and carry the same sign-and-rotation convention as $\hat{B}$ and $\hat{F}_t$, while $\hat{B}\hat{\boldsymbol{\lambda}} = \hat{B}\hat{B}'\widehat{\boldsymbol{\delta}}$ is convention-free; each bootstrap draw recomputes the minimum-distance solution $(\hat{B}^{\star\prime}\hat{B}^{\star})^{-1}\hat{B}^{\star\prime}\hat{\boldsymbol{\delta}}^{\star}$ within the draw (Internet Appendix B.2). When the restriction is not maintained, the same coordinates are read as the loading-span projections of the reduced-form premia. Section 4.2 reports $\hat{\boldsymbol{\lambda}}$ scaled by each factor's estimated daily standard deviation.

Stage 4a (firm-event level γ and exposure Γ). Stage 4a pools the clean-event observations $\mathcal{T}^{(1,0)}$ of (16). This restriction prevents estimation error in the macro-event adjustment from entering the generated factor regressor. It also defines the estimand: macro-date event observations belong to the coincidence regime (c, Γ_a) of (11) and are used by Stage 4b below, so restricting Stage 4a to $\mathcal{T}^{(1,0)}$ keeps its estimand the clean-date pair (γ, Γ). It is n_1, not NT, that governs the Stage-4a rate in Section 3.3. Let $\tilde{R}_{it} := R_{it} - Z'_{i,t-1}\hat{a}_t$ for event observations. Form the regressor vector

$$\widehat{X}_{it} \;:=\; \begin{pmatrix} Z_{i,t-1} \\ Z_{i,t-1}\otimes \hat{F}_t \end{pmatrix} \in \mathbb{R}^{L(K+1)}, \qquad \theta_\gamma^H \;:=\; \begin{pmatrix} \gamma \\ \mathrm{vec}\big((\Gamma H)'\big) \end{pmatrix} \in \mathbb{R}^{L(K+1)},$$

where

$$H \;:=\; (\bar{F}' M_T \hat{F})\,(\hat{F}' M_T \hat{F})^{-1} \tag{22}$$

is the (unobserved) rotation aligning the estimated factor coordinates with the population normalization. The estimated path is stacked in rows like $\bar{F}$ of Assumption 2: $\hat{F} := (\hat{F}_1, \ldots, \hat{F}_T)' \in \mathbb{R}^{T\times K}$, so $\hat{F}' M_T \hat{F}/T$ is the $K \times K$ second-moment matrix of the demeaned estimated factor path and H

is $K \times K$. Stage 4a is the pooled OLS estimator on the event sub-panel:

$$\hat{\theta}_\gamma \;=\; \Bigg(\sum_{(i,t)\in\mathcal{T}^{(1,0)}} \widehat{X}_{it}\, \widehat{X}_{it}' \Bigg)^{-1} \Bigg(\sum_{(i,t)\in\mathcal{T}^{(1,0)}} \widehat{X}_{it}\, \tilde{R}_{it} \Bigg). \tag{23}$$

The first L entries are $\hat{\gamma}$; reshaping the remaining entries yields $\hat{\Gamma}$, which estimates the exposure in the same rotated coordinates as $\hat{F}_t$; the alignment cancels in the pairing, $\hat{\Gamma}\hat{F}_t$ targeting $(\Gamma H)(H^{-1}\bar{F}_t) = \Gamma\bar{F}_t$ free of H. Combined with $\hat{\mu}$ from Stage 3, this delivers the average event effect $z'(\hat{\gamma}+\hat{\Gamma}\hat{\mu})$, invariant to the factor coordinates for the same reason. Because Stage 3 works with calendar-residualized slopes, $\hat{\mu}$ recovers the announcement-free baseline factor mean, so $z'(\hat{\gamma}+\hat{\Gamma}\hat{\mu})$ estimates the average event effect $\tau^\gamma(z)$ of (13).

Multiple event types. Stages 1–3 already accommodate several event types: Stage 1 excludes firms carrying an event of any included type, Stage 2 estimates the M macro premia jointly in (20), and Stage 3 is applied to the slopes residualized on all M calendars. Only Stage 4a must be stacked, because one firm-date observation can belong to several firm-event types at once. In the multivariate framework of Section 2.4, define the type-j block $\widehat{X}_{it}^{(j)} := (Z_{i,t-1}', (Z_{i,t-1}\otimes\hat{F}_t)')' \in \mathbb{R}^{L(K+1)}$ and the stacked Stage-4a design

$$\mathbb{X}_{it} \;:=\; (D_{i,t-1}^{(1)}\,\widehat{X}_{it}^{(1)\prime},\, \ldots,\, D_{i,t-1}^{(J)}\,\widehat{X}_{it}^{(J)\prime})' \;\in \mathbb{R}^{JL(K+1)}.$$

Stacked Stage 4a pools over $\mathcal{T}^{(\bullet)} := \{(i,t) : \max_{j\le J} D_{i,t-1}^{(j)} = 1,\; D_{m,t-1} = 0\}$ and runs OLS of $\tilde{R}_{it} := R_{it} - Z_{i,t-1}'\hat{a}_t$ on $\mathbb{X}_{it}$, returning $\widehat{\boldsymbol{\theta}}_\gamma \in \mathbb{R}^{JL(K+1)}$. Its type-$j$ population block is

$$\theta_{\gamma,j} := \begin{pmatrix} \gamma_j \\ \mathrm{vec}((\Gamma_j H)') \end{pmatrix}.$$

The stacked analogues of the design conditions are part of Assumptions 1 and 3 in Internet Appendix B.2. Coincident events activate several blocks in the same row, so the regression attributes their contributions jointly.

Stage 4b (coincidence level c and exposure Γ_a). The clean and coincidence cells carry different coefficients in (11), so Stage 4a and Stage 4b estimate them by separate regressions on disjoint samples; the split also keeps macro-date returns out of the clean-cell estimand. Let

$$\mathcal{T}^{(1,m)} \;:=\; \{(i,t)\in\mathcal{T}^{(1)} :\; D_{m,t-1} = 1\}$$

collect the firm-event observations on dates carrying at least one macro release (the complement of $\mathcal{T}^{(1,0)}$ within $\mathcal{T}^{(1)}$) and regress $\tilde{R}_{it} = R_{it} - Z'_{i,t-1}\hat{a}_t$ over $\mathcal{T}^{(1,m)}$ on the design

$$\left(D^{(1)}_{m,t-1}Z'_{i,t-1},\ \ldots,\ D^{(M)}_{m,t-1}Z'_{i,t-1},\ (Z_{i,t-1}\otimes \hat{F}_t)'\right)' \in \mathbb{R}^{L(M+K)}.$$

This returns one level block $\hat{c}_g \in \mathbb{R}^L$ per macro type g (the type-g analogue of c in (10), not an entry of it) and one exposure block $\hat{\Gamma}_a$ common to all macro types. Every macro date enters, and a date carrying two releases loads both level blocks additively. Relative to the fully stacked coincidence block in (14), the implementation pools the firm-event types into a single union indicator in both dimensions: $\hat{c}_g$ pools the c_{jg} levels across j, and $\hat{\Gamma}_a$ pools the $\Gamma_{a,j}$ exposures. These are therefore common pooled linear-projection coefficients, not unweighted averages of the pair-specific structural coefficients. Their weights depend on type-specific regressor second moments and overlap as well as on event frequencies; only under common second moments and no overlap do they reduce to event-frequency-weighted averages; the coincidence cell is too thin to support a separate exposure block for each of the JM pairs. The estimated $(\hat{c}_g, \hat{\Gamma}_a)$ is therefore a reduced-form description of coincidence-date conditional means in the sense of Proposition 1(iv), not a structural announcement effect.

The estimates $(\widehat{\boldsymbol{\delta}}, \hat{\alpha}, \hat{B}, \hat{\gamma}, \hat{\Gamma})$ do not depend on the block. From the block we report the structural announcement-day level, $\hat{\gamma}_{a,g} := \widehat{c}_g - \widehat{\Gamma}_a\hat{\lambda}_g$, the sample analogue of Proposition 1(v). Because $\widehat{\Gamma}_a$ is pooled, the subtraction recovers a type-g structural level under a common announcement-day exposure across firm-event types, a restriction we maintain for this block; without it, $\hat{\gamma}_{a,g}$ is read as the pooled announcement-day level. We report it together with $\widehat{\Gamma}_a$ and the announcement-day average event effect $z'(\widehat{c}_g + \widehat{\Gamma}_a\widehat{\mu})$, the sample counterpart of the third contrast $\tau^{\gamma\delta}$ of Section 2.1.

Stages 1–3 are the regressed-PCA estimator of CRW (their eq. 10) applied to the non-event sub-panel after the time-series projection on the macro calendar. Stage 4a is a generated-regressor pooled OLS in the spirit of Pagan (1984), Murphy and Topel (1985), and Bai and Ng (2006), with $\hat{F}_t$ entering as a generated regressor whose pre-rotation estimation error is governed by Theorem 1 in Internet Appendix B.2 (summarized in Section 3.3); the asset-pricing antecedent is the errors-in-variables correction of Shanken (1992) for estimated betas in two-pass regressions.

3.3 Large-Sample Properties and Inference

The four stages are estimated sequentially but analyzed jointly. The formal results are stated in Appendix B.2 and proved in Appendix C. Under Assumptions 1–3, the four-stage estimates of $(\boldsymbol{\delta}, \alpha, B, \gamma, \Gamma)$ are consistent. The loading and exposure blocks are expressed in the coordinates aligned by H in (22).

Under Assumption 4 and $T/N \to 0$, Theorem 1 establishes joint asymptotic normality at mixed rates. The macro premia are estimated from calendar variation and converge at rate $\sqrt{T}$. The baseline and loading blocks use both dimensions of the panel and converge at rate $\sqrt{NT}$. The

firm-event blocks converge at rate $\sqrt{n_1}$. The aspect-ratio condition is weakened for the jackknifed estimator by Proposition 2. The theorem also covers a fixed number of firm-event and macro-event types through stacking.

The limiting covariance matrix is singular because the $\sqrt{T}$-scaled macro block is supported on $\mathrm{col}(B)$ and $\alpha' B = 0$ holds in sample. Inference is therefore stated for scalar functionals not the full parameter vector. Under Assumption 5, Theorem 2 establishes conditional validity of the split firm and date multiplier bootstrap. Corollaries 1 and 2 provide inference for the rotation-invariant functionals and average event effects reported in the paper. Appendix A examines the finite-sample performance of the approximation. Stage 4a is not an ordinary final-step regression as its benchmark and factor regressors are estimated, and their errors are shared by every event firm on a date. The relevant covariance is therefore the joint one of all four stages, with cross-stage terms that conventional Stage-4a OLS standard errors omit (Appendix B constructs it blockwise).

Bootstrap inference. Every standard error, interval, and t-ratio for the model parameters reported in Section 4 is computed from a split firm/date multiplier bootstrap (Appendix B.3), not from a plug-in estimate of $\mathcal{V}$; the portfolio exhibits use conventional time-series statistics on realized returns, identified as such where they appear: with mixed rates, a singular covariance, and cross-stage generated-regressor terms, direct studentization is impractical, whereas the bootstrap estimates the law of the four-stage estimates directly. One draw proceeds as follows.

(1) Draw i.i.d. mean-one firm weights $\{w_i^\star\}$, shared by all observations of firm i, and independent mean-zero date signs $\{w_t^\star\}$; the implementation uses standard exponentials and Rademacher signs.

(2) Stage 1*: rerun the control regressions with the firm weights,

$$\widehat{a}_t^\star = \left(\textstyle\sum_{i\in\mathcal{T}_t^{(0)}} w_i^\star Z_{i,t-1} Z_{i,t-1}'\right)^{-1} \sum_{i\in\mathcal{T}_t^{(0)}} w_i^\star Z_{i,t-1} R_{it}.$$

(3) Stage 2*: with $\widehat{v}_t := \widehat{a}_t^0 - \bar{\widehat{a}}^0$ the demeaned Stage-2 residual, the reported macro draw $\hat{\boldsymbol{\delta}}^\star$ is the calendar regression (20) applied to $\widehat{a}_t^\star + w_t^\star \hat{B}\hat{B}'\widehat{v}_t$. The input passed to Stage 3 is residualized with the firm-only coefficient $\hat{\boldsymbol{\delta}}^{\star,N}$, the same regression without the date-wild term, as $\widehat{a}_t^{0,\star} = \widehat{a}_t^\star - \hat{\boldsymbol{\delta}}^{\star,N} \boldsymbol{D}_{t-1}^m$, so the exact calendar annihilation carries over to the draw.

(4) Stage 3*: PCA on $\{\widehat{a}_t^{0,\star}\}$ gives $(\hat{B}^\star, \hat{\alpha}^\star, \hat{F}_t^\star)$; loading and exposure draws are centered at $\hat{B}H^\star$ and $\hat{\Gamma}H^\star$, with the regression alignment $H^\star := (\hat{F}'M_T\hat{F}^\star)(\hat{F}^{\star\prime}M_T\hat{F}^\star)^{-1}$ and column signs matched; the factor announcement premium draw is $\hat{\boldsymbol{\lambda}}^\star = (\hat{B}^{\star\prime}\hat{B}^\star)^{-1}\hat{B}^{\star\prime}\hat{\boldsymbol{\delta}}^\star$.

(5) Stages 4a*–4b*: $w_i^\star$-weighted OLS of $R_{it} - Z_{i,t-1}'\widehat{a}_t^\star$ on each cell's own design with $\hat{F}_t^\star$ in place of $\hat{F}_t$: the clean-cell design $(Z_{i,t-1}', (Z_{i,t-1}\otimes\hat{F}_t^\star)')'$ on $\mathcal{T}^{(1,0)}$ and the Stage-4b coincidence design on $\mathcal{T}^{(1,m)}$, with every reported statistic recomputed within the draw. For the average event effects, the date signs $w_t^\star$ of step (1) also redraw the clean- and announcement-calendar

means of the draw's factor path, so the draw dispersion prices the sampling variation of the factor-path means around their population values (Corollary 2).

Standard errors are standard deviations across 200 draws 500 for the pooled specification; percentile and studentized intervals use the same draws. Theorem 2 and Corollaries 1 and 2 are what license these numbers, with Proposition 3 covering the coincidence pair. The draw-by-draw algebra, including why the residualization must use the firm-only Stage-2 coefficient, is in Internet Appendix B.3.

Jackknife bias correction. The exposure and loading entries are consistently estimated; when T is large relative to N, their largest components carry a finite-sample attenuation, documented in Internet Appendix A. To remove it, the estimates are combined with a cross-sectional half-panel jackknife, which draws S independent random partitions of the firms into halves, re-runs the four stages on each half (each half's loading and exposure blocks mapped into the full-sample estimated coordinates by regressing its recovered factor path on the full-sample path, an operation that uses no unobserved quantity) and combines

$$\widehat{\vartheta}^{\mathrm{jk}} := 2\,\widehat{\vartheta} - \frac{1}{2S}\sum_{s=1}^{S}\big(\widehat{\vartheta}_{(s,1)} + \widehat{\vartheta}_{(s,2)}\big).$$

The generated-regressor bias of order N^{-1} doubles on a half cross-section, so the combination removes it at first order while leaving the first-order influence, and with it every reported bootstrap standard error, unchanged. Proposition 2 provides the corrected estimator's centered limit theory, under an explicit stability condition on the half-sample bias expansions, at the much weaker sample-size condition $T/N^2 \to 0$, which the panel satisfies. The bias the correction removes is material in the loading and exposure blocks; the loading block is not separately reported, and the reported quantities built from it, the minimum-distance factor premia and the span decomposition, are essentially unchanged under the corrected loadings, while for the level blocks the correction trades a negligible bias for half-sample noise, which averaging over the S splits suppresses. The empirical exposure estimates are reported with the correction applied (Section 4.2).

Stage 2 estimates $\boldsymbol{\delta}$ as free L-vectors, so the factor-channel implication $(I_L - BB')\delta = 0$ of (10) is not imposed anywhere in estimation. Section 4.2 reports the descriptive span decomposition of the estimated premia. Under either reading of the decomposition, $\hat{B}'\hat{\delta}$ carries the reduced-form interpretation as the loading-span coordinate of the premium, the event-calendar analogue of the projection by which Giglio and Xiu (2021) recover an observable factor's premium, while its structural reading as the factor announcement premium, and the other structural quantities of Proposition 1(v), are reported only under the explicitly maintained factor-channel restriction (the minimum-distance $\hat{\boldsymbol{\lambda}}$ of Section 3.2); every other estimate in Section 4 is a reduced-form conditional-return coefficient.

4 Empirical Analysis

We apply the model to U.S. common stocks from 1980 through 2024. We first describe the return panel, event calendars, and firm characteristics. We then estimate the pooled and type-specific event components and examine exposure to industry peer announcements. The remaining exercises evaluate the fitted signals across held-out firms and in portfolios formed from recorded event observations.

4.1 Data and Event Description

The return panel uses daily CRSP returns from 1980 through 2024. We retain domestic common shares listed on the NYSE, AMEX, or Nasdaq and require a lagged closing price of at least \$5. Observations with missing returns are excluded, and excess returns are measured relative to the daily risk-free rate. The final panel contains 42.3 million firm-day observations for about $21,000$ firms over $11,323$ trading dates. The average daily cross-section contains approximately $3,700$ stocks.

Firm-specific events consist of earnings announcements and distribution ex-dates. Earnings dates come from Compustat quarterly fundamentals. [1] Distribution dates come from the CRSP distribution file. The broad indicator includes every distribution ex-date, while a separate specification retains cash distributions identified by CRSP codes 12xx. The firm-event indicator equals one when either event occurs.

Macro events are scheduled CPI, PPI, Employment Situation, and FOMC announcements. The macro-event indicator equals one when at least one of these releases occurs. Of the $11,323$ sample dates, $9,433$ are release-free and $1,890$ contain a macro release. Only 61 dates contain two releases, most often CPI or PPI together with an FOMC decision, and none contains three. Firm events account for 2.2% of firm-day observations. In the mutually exclusive classification, 81.4% of observations contain no event, 16.4% contain only a macro event, 1.9% contain only a firm event, and 0.3% contain both.

Figure 1 shows the event structure in 2024. The filtered panel contains $3,428$ firms and 252 trading days, with $10,696$ earnings-announcement firm-dates and $5,266$ distribution ex-dates. The macro calendar contains 8 scheduled FOMC dates and 12 dates for each of the Employment, CPI, and PPI releases.

The figure also shows the two sources of variation used in estimation. Changes in date-level return slopes between macro-event and other dates identify $Z'_{i,t-1}\delta$. Comparisons between firms with and without an event on the same date identify $Z'_{i,t-1}\gamma + Z'_{i,t-1}\Gamma\bar{F}_t$, conditional on the common factor realization. Macro events therefore provide variation across dates, while firm events provide

[1] Following the recurring-event literature (Hartzmark and Solomon, 2018), we condition on the recorded daily event date. Announcement dates are also firm-chosen because firms can advance or delay reports (deHaan et al., 2015; Johnson and So, 2018).

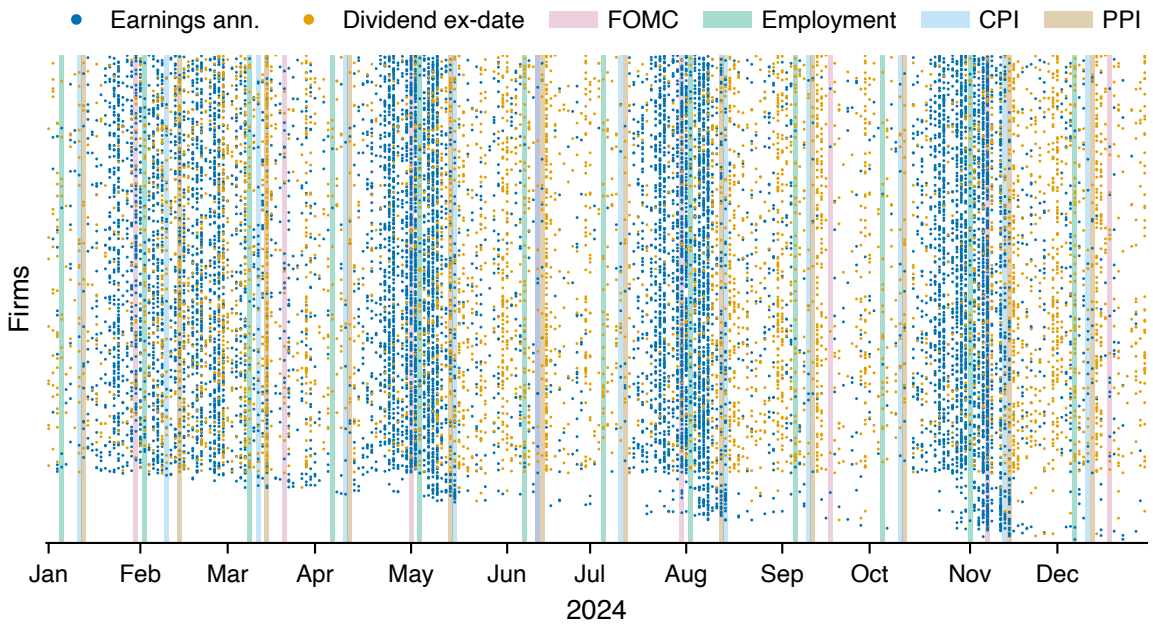


Figure 1: Example event calendar for 2024. Rows are firms and columns are trading days. Colored background columns mark macro-event dates for FOMC, Employment, CPI, or PPI announcements, while dots mark earnings announcements and distribution ex-dates.

variation across firms within a date.

Figure 2 reports unconditional average excess returns by event type. Firm-event observations earn 20.1 basis points per day, compared with 7.6 basis points for macro-event observations and 1.9 basis points for observations with neither event. Distribution ex-dates have the largest type-specific mean at 36.2 basis points, followed by FOMC dates at 18 and earnings announcements at 11. Employment, PPI, and CPI dates have means of 10.2, 5, and 3.5 basis points. These are unconditional comparisons. The estimates below condition on characteristics and separate the common factor component.

Table 1 separates cash distributions from other CRSP distribution records. The mean excess return is 36.21 basis points under the broad definition and 29.98 basis points for cash distributions. A small group of noncash-only observations has a mean of 137.30 basis points and accounts for much of the difference. On cash-distribution dates, mean CRSP total return including distributions, `RET`, is 31.76 basis points, while mean return excluding distributions, `RETX`, is -48.91 basis points. Their difference of 80.67 basis points is the contribution of the distribution to the holding-period return.

We compute the price-drop ratio from the CRSP distribution and daily stock files because it requires both the distribution amount and the lagged price. The sample contains 720,380 cash-distribution firm–ex-dates from 1980 through 2024. We retain 618,126 observations with a lagged price above \$1 and a ratio between -5 and 5. The ratio of the close-to-close price decline to the distribution amount rises from 0.53 in the 1980s to 0.66 in the 2000s and 0.79 in 2020–2024. Over

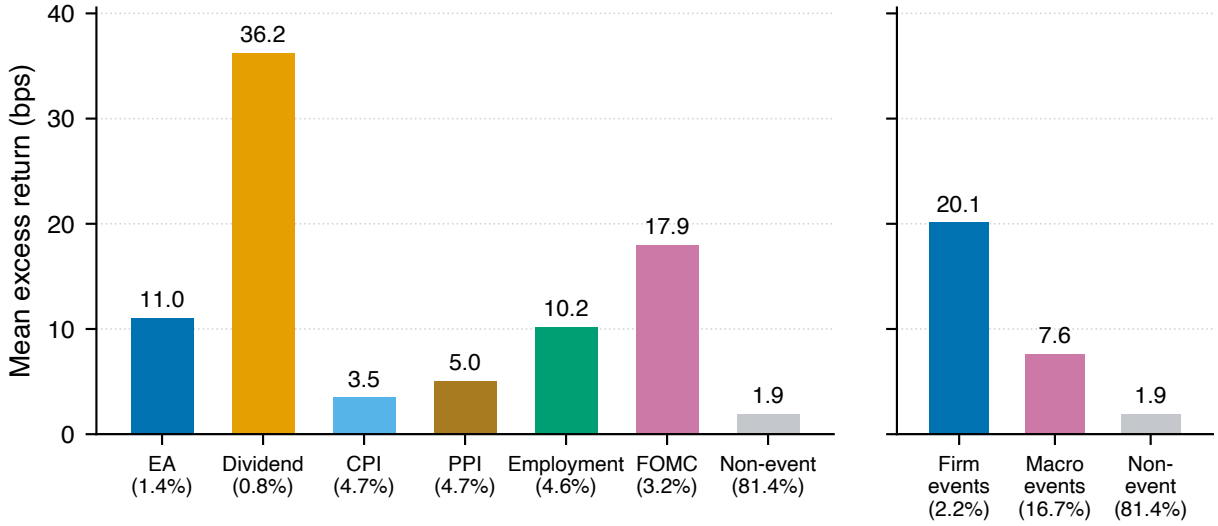


Figure 2: Average daily excess returns by event type and by broad event class. The event-type panel counts an observation in each event category whose indicator equals one; the event-class panel aggregates firm events and macro events separately. Event probabilities are shown in parentheses on the x-axis.

Table 1: CRSP distribution ex-date returns

Distribution definition	Firm–ex-dates	Mean excess return (bps)
All CRSP distribution records	334,958	36.21
Cash distributions (codes 12xx)	315,497	29.98
Code 1232	298,464	28.88
Noncash-only records	19,461	137.30

Notes: The table matches the constructed distribution-event indicator to the CRSP distribution file after applying the return-panel filters. All firm–ex-dates are matched. Because a firm can have more than one distribution record on the same date, the noncash-only row includes firm–ex-dates with no distribution carrying a 12xx code.

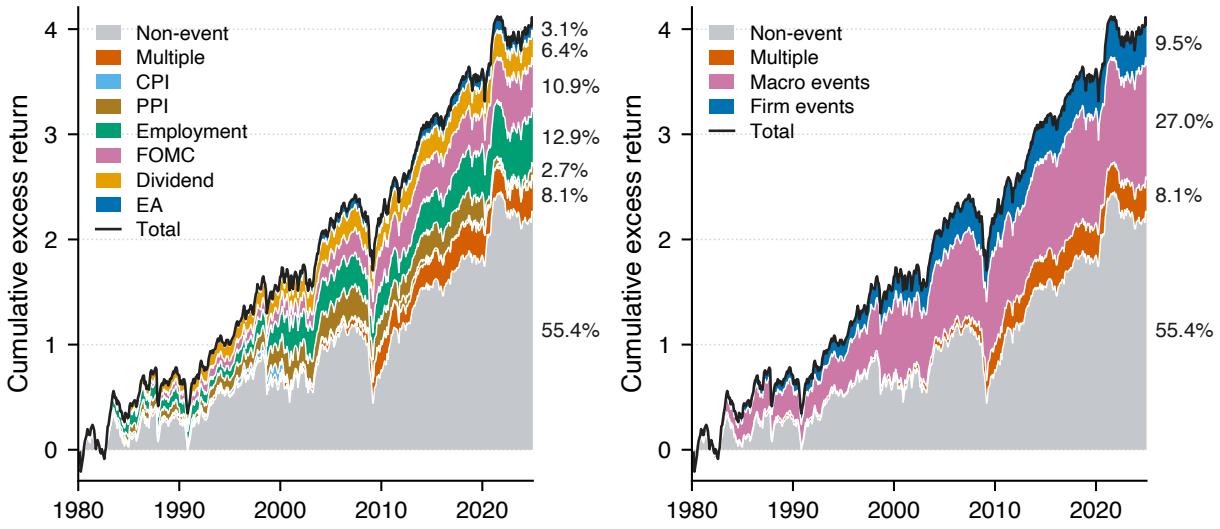


Figure 3: Cumulative equally weighted excess return decomposed by event category. Each layer is the cumulative contribution of the corresponding event category. The right panel aggregates categories into firm events, macro events, multiple-event observations, and non-events.

the same periods, the mean raw ex-date return falls from 48 to 8 basis points. Excess returns in the filtered panel show the same decline, from 38.79 to 4.86 basis points. Their levels differ because the panel applies additional filters and subtracts the risk-free rate.

The price-drop ratio also varies with the tax and trading environment. It rises from 0.52 in 1980–1986 to 0.60 in 1987–1997 and from 0.56 in 2001–2002 to 0.70 in 2003–2012. Before 2001, mean ex-date returns range from 21 to 59 basis points across dividend-yield quintiles. After 2001, the range narrows to 11–19 basis points. Bid–ask bounce can also inflate ex-date returns measured with close-to-close prices (Frank and Jagannathan, 1998). In every decade, the return implied by the incomplete price drop matches or exceeds the realized mean ex-date return.

These patterns match the ex-dividend-day evidence in Elton and Gruber (1970), Kalay (1982), Eades et al. (1984), Frank and Jagannathan (1998), and Graham et al. (2003). Aggregating the pattern over the dividend calendar is related to the dividend-month premium of Hartzmark and Solomon (2013). We interpret the cash-distribution estimate as the predictable return associated with the classical incomplete ex-dividend price adjustment. The ex-date is a prescheduled firm event that implements an earlier payout decision. Recovering this established pattern provides a benchmark for the estimator.

The cumulative decomposition in Figure 3 shows that event observations contribute more to realized returns than their frequency suggests. Observations with no event account for 55.4% of cumulative equally weighted excess returns. The four macro-only categories contribute 27.0%, led by Employment and FOMC dates. Firm-only categories contribute 9.5%, and observations with multiple event indicators contribute 8.1%.

Firm characteristics come from the Jensen–Kelly–Pedersen library.[2] Monthly observations are assigned to daily returns using information available at the end of the preceding month. Missing values are replaced by the cross-sectional median within date. Each characteristic is then rank-normalized to $[-1, 1]$. The specification uses 29 characteristics and an intercept, giving $L = 30$.

The characteristics cover profitability, investment, financing, valuation, payout, risk, return dynamics, and trading conditions. Profitability measures include profit margin and returns on assets, equity, investment, and physical capital. Investment and financing variables include asset and profit growth, investment, depreciation growth, issuance, equity payout, leverage, financial soundness, and solvency. Valuation and payout measures include book-to-market, profit-to-market, equity payout-to-market, dividend-to-market, and size. The remaining variables are short-term reversal, long-term momentum, seasonality, idiosyncratic and total return volatility, CAPM beta, zero-trade days, turnover, liquidity, and activity. Internet Appendix D gives the definitions and update frequencies.

4.2 Estimation Results

We first report the pooled estimates and their cross-sectional heterogeneity. We then examine differences across event types, coincident firm and macro events, and exposure to industry peer announcements.

4.2.1 Pooled Estimates

The estimation sample contains 42.3 million firm-day observations, about 21,000 firms, and 11,323 trading dates. We estimate both the pooled specification and the stacked specification of Section 2.4 using $K = 3$ latent factors.[3] The pooled specification combines earnings announcements and distribution ex-dates in one firm-event indicator and combines CPI, PPI, Employment, and FOMC announcements in one macro-event indicator. The stacked specification estimates separate coefficients for each type and allows the indicators to overlap. Because the characteristics are rank-normalized and centered within date, each intercept is the conditional return at the cross-sectional mean of the characteristics. Standard errors use the corrected split-weighted bootstrap of Internet Appendix B.3, with PCA re-estimated in every draw.[4]

Figure 4 reports the pooled coefficient vectors. The event-free intercept is 0.7 basis points per day. The pooled macro-event component is 4.8 basis points and is less precisely estimated. The pooled firm-event component is 18.2 basis points, about 27 times the event-free intercept and close to the raw firm-event mean reported in Section 4.1. Firm-event dates therefore carry substan-

[2] https://jkpfactors.com/

[3] The first factor is dominant, and three factors explain 60% of the macro-residualized slope variation. Re-estimation with $K \in \{2, 4, 5\}$ leaves the earnings, distribution, and FOMC intercepts nearly unchanged. The event-free alpha and factor-dependent objects are more sensitive to K.

[4] Inference uses 500 bootstrap draws for the pooled specification and 200 for the stacked specifications. The coefficient figures report pointwise significance without adjustment for multiple testing.

tially higher conditional returns than ordinary trading dates even before using characteristics to distinguish among firms.

The three intercepts describe different comparisons. The event-free intercept is the fitted return when neither event indicator is active. The macro coefficient measures the change in date-level return slopes on macro-announcement dates. The firm-event coefficient compares firms with and without a firm event on the same nonmacro date, conditional on the common factor realization. The pooled firm-event estimate is a descriptive average of earnings announcements and distribution ex-dates, while the pooled macro estimate averages four different announcement calendars. We separate these components below. The slopes show that event status also changes the cross-section of returns. In the event-free block, long-term momentum is the strongest positive coefficient. Dividend-to-market, book-to-market, and several profitability measures are also positive. Short-term reversal and idiosyncratic volatility have the largest negative coefficients, with additional negative relations for issuance, turnover, investment, and asset growth. This block resembles the familiar non-event return cross-section in which recent winners, value firms, and profitable firms earn higher fitted returns.

The macro-event block has a different profile. CAPM beta, idiosyncratic volatility, size, long-term momentum, and issuance have positive slopes, while profit-to-market and dividend-to-market have negative slopes. The positive beta coefficient is particularly important because scheduled macro announcements are dates on which investors learn about aggregate conditions (Savor and Wilson, 2014). The volatility coefficients also indicate that the macro component is not a common shift shared equally by all firms. It is larger for firms whose characteristics imply greater sensitivity to aggregate news.

The firm-event block differs from both. Its largest positive coefficients include long-term momentum, dividend-to-market, leverage, zero-trade days, and return on physical capital. Its largest negative coefficients include equity payout-to-market, turnover, and size. Momentum links the event return to prior price performance, while the size, turnover, liquidity, and zero-trade-day coefficients point to an important role for the trading environment around firm events. These relations are consistent with price impact or limited-arbitrage amplification, although they do not separately identify those channels. The negative turnover coefficient refers to average share turnover over the prior six months and therefore does not contradict evidence based on volume in earlier announcement windows (Frazzini and Lamont, 2007).

The contrast across the three blocks is important for interpretation. A characteristic can predict ordinary returns without having the same relation to the additional return on an event date. Likewise, a characteristic can have different slopes in the macro and firm-event blocks because the two events supply different comparisons. The estimates therefore describe changes in the cross-sectional organization of returns, not only changes in their average level. The coefficient signs alone cannot distinguish risk, attention, liquidity, and price-pressure explanations.

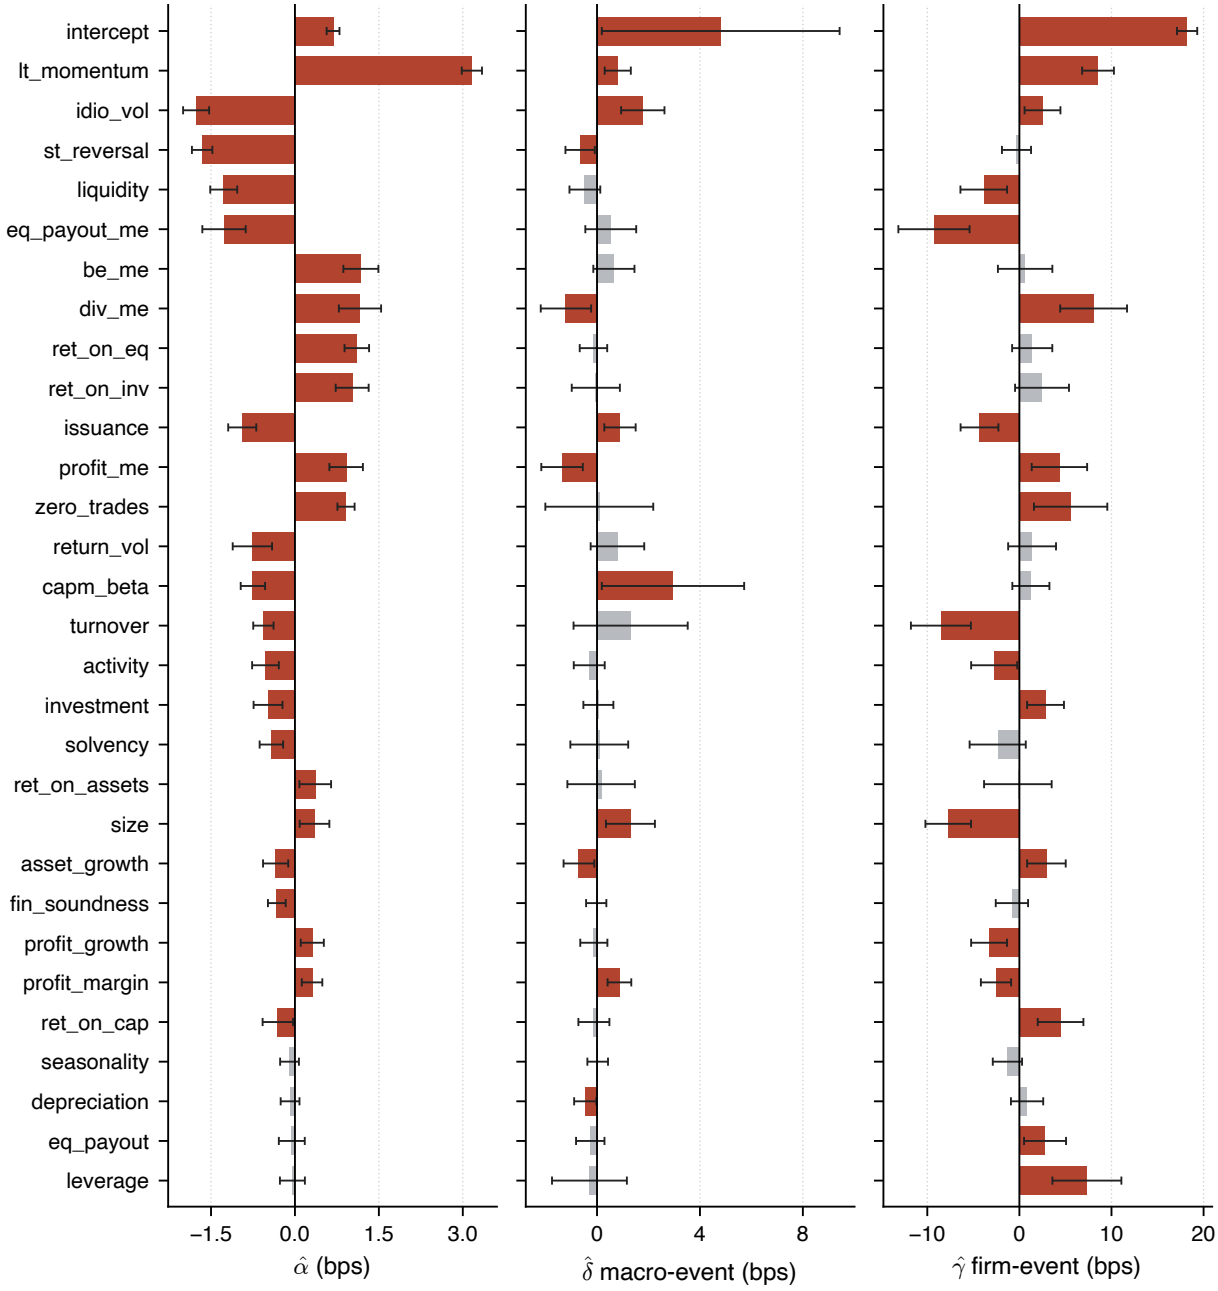


Figure 4: Estimated coefficient vectors. The panels report the event-free alpha $\widehat{\alpha}$, the macro-event slope $\widehat{\delta}$, and the firm-event level $\widehat{\gamma}$. Coefficients are in basis points. Highlighted bars denote statistical significance at the 10% level using the corrected split-weighted bootstrap with PCA re-estimated in each of 200 draws.

Table 2: Average conditional-return components by event type

Component	Estimate (bps)	Bootstrap SE	t-statistic
Event-free baseline	0.68***	0.07	9.78
Firm events, pooled	18.21***	0.67	27.22
Earnings	9.51***	0.79	12.06
Distribution ex-date, broad	51.96***	1.18	44.14
Cash distribution (12xx)	33.16***	1.00	33.02
Macro events, pooled	4.81†	2.81	1.71
CPI	0.54	4.64	0.12
PPI	0.40	5.15	0.08
Employment	7.56	5.03	1.50
FOMC	17.01***	6.23	2.73

Notes: The table reports the intercept from each estimated coefficient vector, in basis points, with bootstrap standard errors and t-statistics. Characteristics are rank-normalized within date with average ranks for ties, so each has cross-sectional mean zero and the intercept evaluates the component at the characteristic mean. Standard errors use the corrected split-weighted bootstrap with PCA re-estimated in every draw: 500 draws for the pooled rows and the event-free baseline, 200 draws for the event-type rows. *, **, and *** denote two-sided pointwise significance at the 10%, 5%, and 1% levels. †: the pooled macro t of 1.71 clears the 10% threshold mechanically but falls below it once the documented worst-case bootstrap understatement of the macro block (z-variances up to 1.42) is applied; we do not mark it as significant. All rows except "Cash distribution (12xx)" use the broad CRSP distribution indicator. The cash row comes from a complete four-stage re-estimation after replacing that indicator with 12xx firm–ex-dates.

4.2.2 Event-Type Decomposition

The pooled estimates combine events with different average returns. Table 2 reports the type-specific intercepts. The earnings-announcement estimate is 9.51 basis points.[5] The estimate for all CRSP distribution records is 51.96 basis points, while a re-estimation restricted to cash distributions gives 33.16 basis points. As shown in Section 4.1, the broad category includes a small set of high-return noncash records, and the cash result follows the classical incomplete ex-dividend price adjustment. We therefore use the cash estimate when interpreting the distribution result.

The pooled macro estimate also masks substantial differences across announcement types. The premium is concentrated on FOMC dates, where the estimate is 17.01 basis points. CPI and PPI estimates are close to zero, and the Employment estimate is positive but less significant. The FOMC estimate is of the same order as the announcement-window evidence in Savor and Wilson (2014) and Lucca and Moench (2015).

The type-specific slopes confirm that pooling conceals different cross-sectional patterns. Earnings returns are strongly related to momentum and trading characteristics. Distribution returns vary more with size, payout, volatility, and microstructure characteristics. The macro slopes are also heterogeneous. The positive beta relation is concentrated on FOMC and Employment dates, while CPI and PPI display weaker patterns. Figures 5 and 6 report the complete coefficient vectors.

[5] Value weighting lowers the earnings estimate to 7.19 basis points but leaves the distribution and FOMC estimates nearly unchanged. Extending the earnings indicator through the next trading day adds little to the estimated premium. Delayed announcements have negative returns, while on-time and early announcements have positive returns. Compustat and I/B/E/S announcement dates generally agree where both are available, especially after 2000.

We next decompose each macro coefficient into its projection onto the estimated loading space and its orthogonal component. Stage 2 estimates the macro coefficients without imposing a factor-channel restriction, so this decomposition is descriptive. For the FOMC coefficient, 82% of its squared norm lies in col($\widehat{B}$). The corresponding share is 69% for Employment and is small for the near-zero CPI and PPI coefficients. At the cross-sectional mean, the FOMC span component is 17.6 basis points and its orthogonal component is -0.6 basis points.

Under the additional restriction $\delta_g = B\lambda_g$, the loading-space coordinates can be read as announcement-day shifts in factor means. On this reading, the FOMC shift is concentrated in the first factor and equals 17.3% of its estimated daily standard deviation, with a bootstrap standard error of 6.0%. Without the restriction, the same estimate remains the loading-space coordinate of the reduced-form FOMC coefficient. The restriction is not imposed in estimation, and the individual coordinates are conditional on the Stage-3 factor convention.

Figure 7 reports the firm-event factor-exposure matrix $\widehat{\Gamma}$. Whereas $\widehat{\gamma}$ describes the level of the firm-event return, $\widehat{\Gamma}$ describes how that return changes with the realized factors. We report the exposure block with the half-panel jackknife of Proposition 2. The correction is material for some individual entries, especially in the third factor, but has little effect on the aggregate decomposition because the realized factor path averages near zero. Individual entries are read under the Stage-3 factor orientation.

We then examine firm events that occur on macro-announcement dates. Table 3 reports two estimates for each macro type. The level $\hat{\gamma}_{a,g}$ isolates the firm-event component that is not associated with factor exposure. The average effect $\hat{\tau}_g^{\gamma\delta}$ adds the compensation associated with the announcement-day exposure, evaluated at the mean factor level for macro type g. Both quantities are evaluated at the cross-sectional mean of the characteristics.

The average firm-event effect is 19.97 basis points on CPI dates, 8.96 basis points on PPI dates, 9.99 basis points on Employment dates, and 14.64 basis points on FOMC dates. The corresponding event levels are similar, which indicates that the factor-exposure component is small at the cross-sectional mean. The estimates remain positive across all four macro calendars, although their magnitudes differ substantially.

Figure 8 reports the characteristic slopes in the coincidence cells. Their broad organization resembles the clean firm-event slopes, with positive value and momentum coordinates and negative size, issuance, and payout coordinates. The smaller cells make individual slopes less precise. The clearest departures occur on FOMC dates, where the premium tilts toward profitability, and on CPI dates, where it tilts toward book-to-market and away from size and equity payout-to-market.

Figure 9 summarizes the fitted contribution of the separate event types. Earnings and distribution ex-dates make distinct contributions to the firm-event component, while the macro contribution is concentrated on FOMC dates. The figure uses the stacked event-type specification.

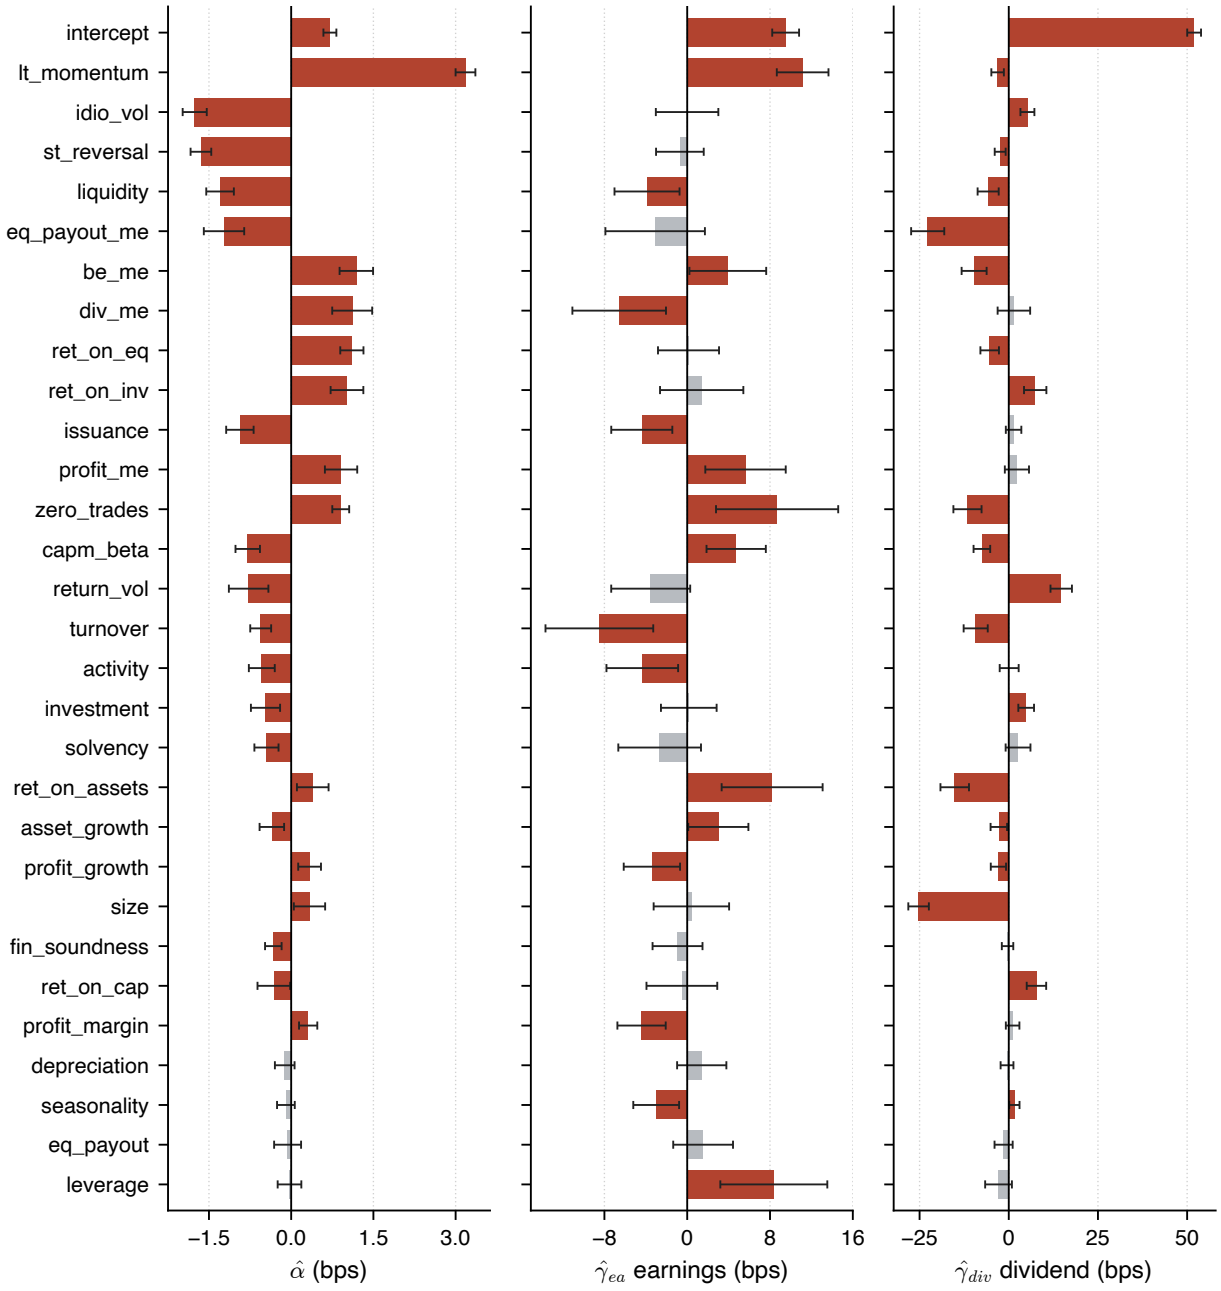


Figure 5: Event-free and firm-event coefficient estimates from the stacked specification. The panels report $\widehat{\alpha}$, $\widehat{\gamma}_{\text{ea}}$, and $\widehat{\gamma}_{\text{dist}}$ in basis points. The plotting code's "dividend" label denotes the broad CRSP distribution ex-date event defined in Section 4.1. Highlighted bars denote pointwise significance at the 10% level under the corrected split-weighted bootstrap.

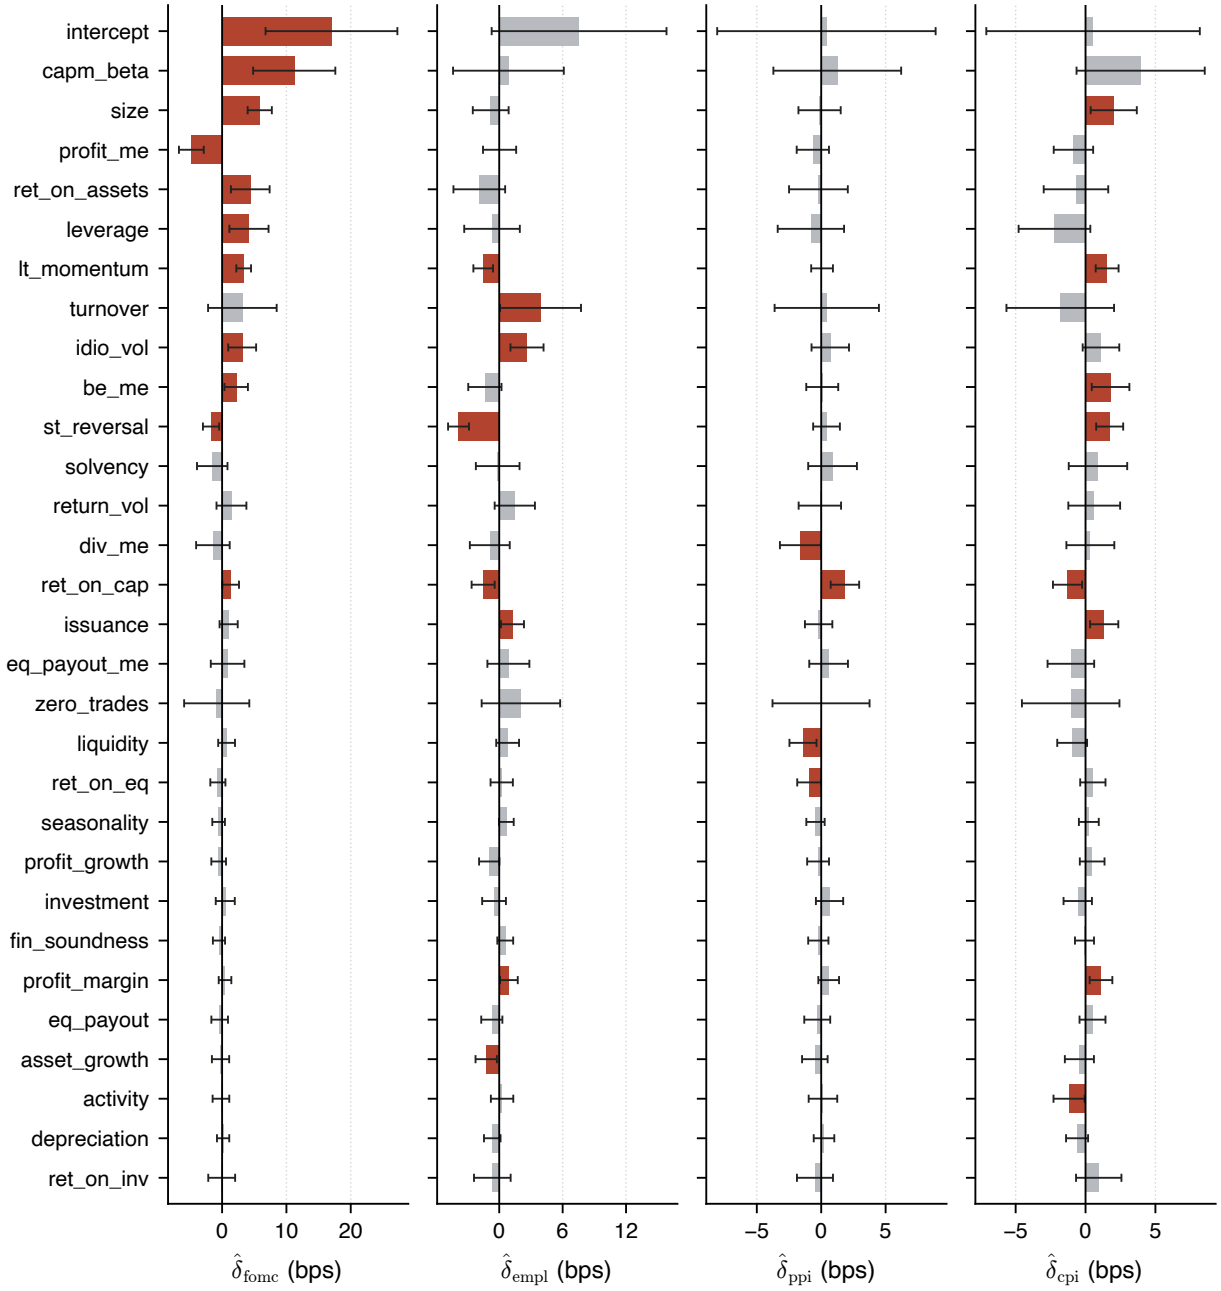


Figure 6: Macro-event coefficient estimates from the stacked specification. The panels report the FOMC, Employment, PPI, and CPI coefficient vectors in basis points. Highlighted bars denote pointwise significance at the 10% level under the corrected split-weighted bootstrap with PCA re-estimated in each of 200 draws.

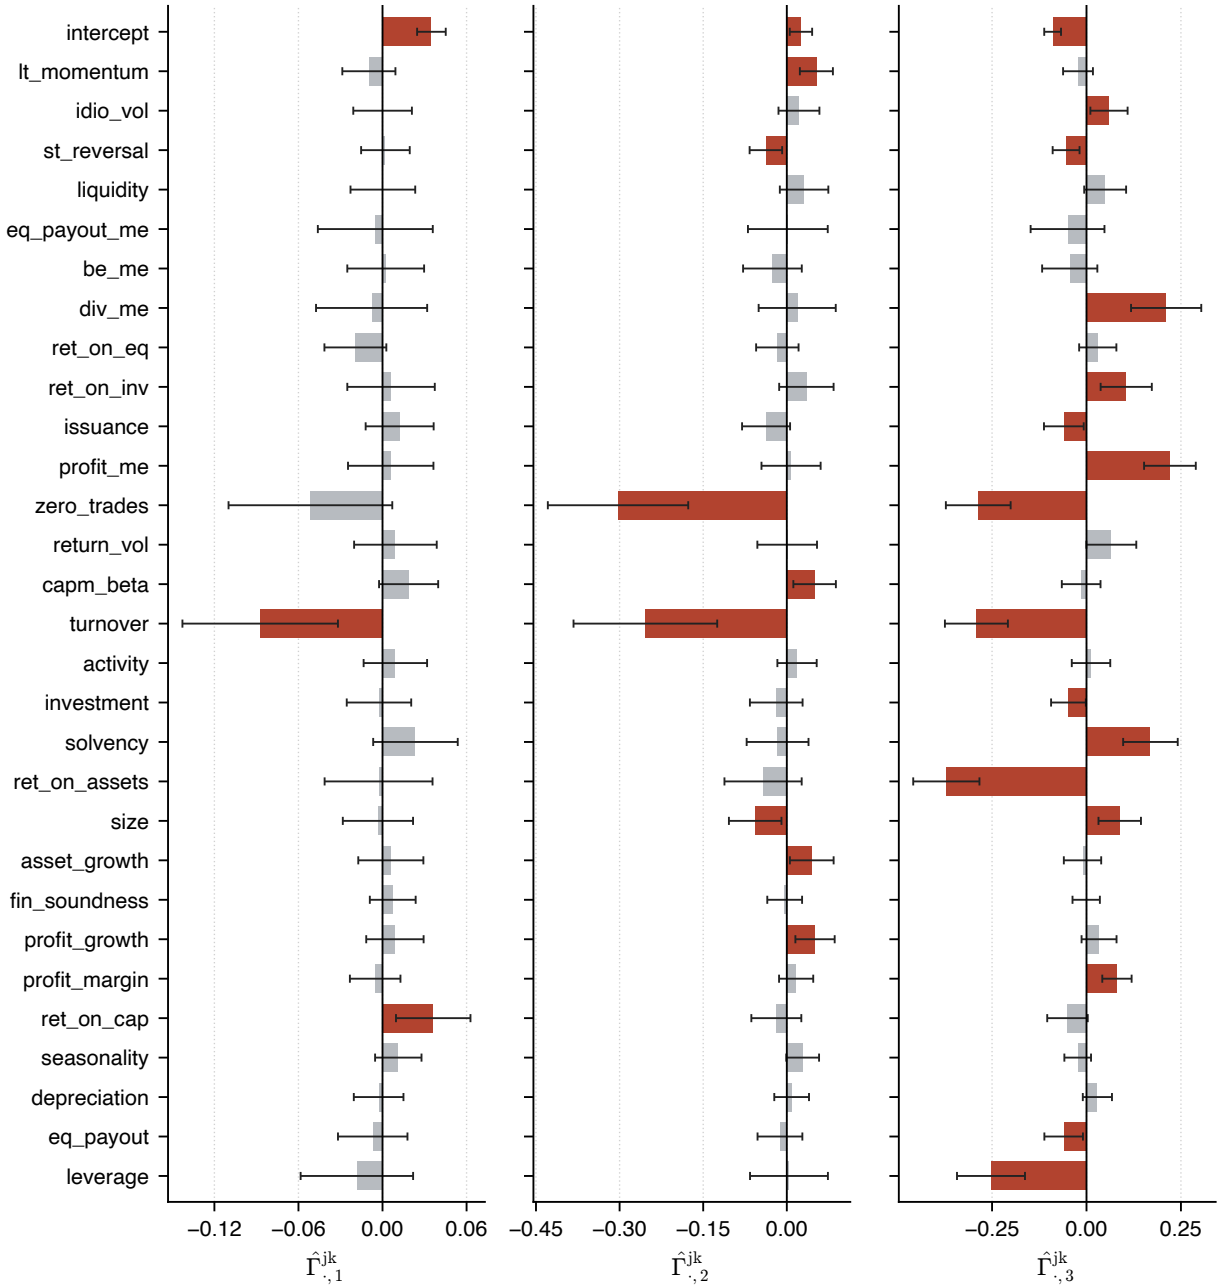


Figure 7: Estimated firm-event factor exposure $\widehat{\Gamma}$ for $K = 3$ latent factors, with the half-panel jackknife of Proposition 2 applied ($S = 100$ splits). Each panel reports one factor column. Highlighted bars denote pointwise significance at the 10% level, with intervals centered at the jackknife-corrected estimates.

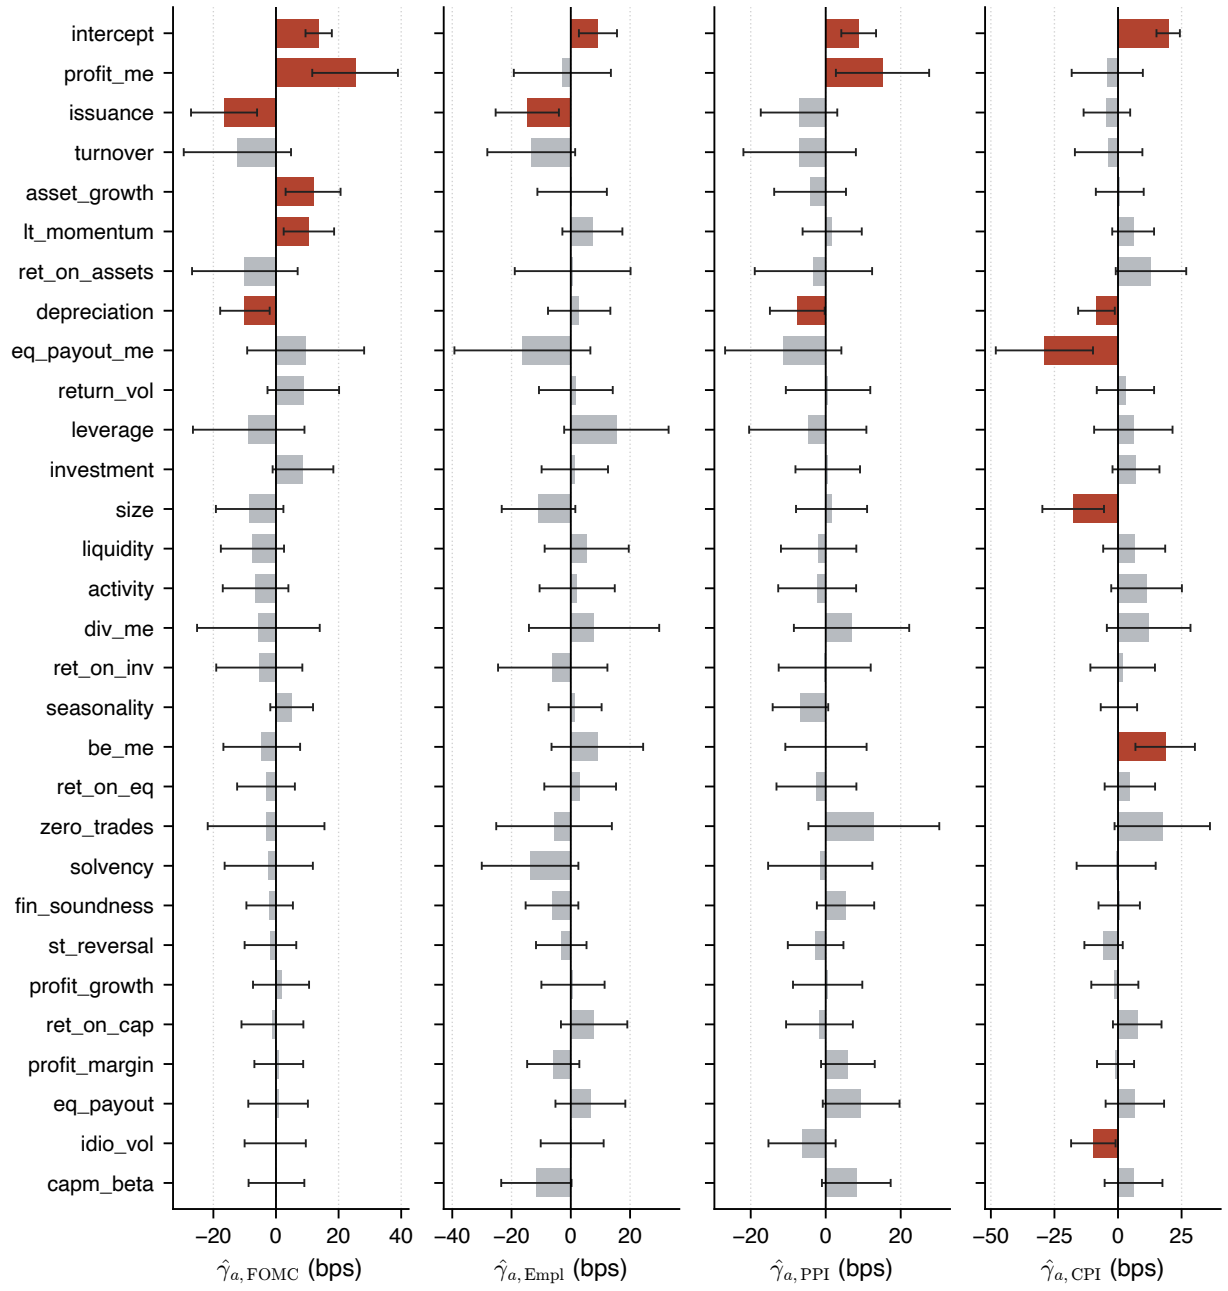


Figure 8: Announcement-day firm-event level: characteristic slopes. The panels report the coefficient vectors $\hat{\gamma}_a$ for the FOMC, Employment, PPI, and CPI coincidence cells in basis points, from the Stage-4a/4b system with standard errors from 200 joint bootstrap draws. Highlighted bars denote pointwise significance at the 10% level; γ_a is rotation invariant, so the marker convention of the invariant exhibits applies. Ordering follows the FOMC panel.

Table 3: Firm-event effects on macro-announcement dates

Macro type	$\hat{\gamma}_{a,g}$	$\hat{\tau}_g^{\gamma\delta}$
CPI	19.73 (2.80)	19.97 (2.92)
PPI	8.79 (2.82)	8.96 (2.96)
Employment	9.16 (3.91)	9.99 (3.89)
FOMC	13.64 (2.56)	14.64 (2.79)

Notes: Entries are basis points at the cross-sectional mean of the rank-normalized characteristics, with bootstrap standard errors in parentheses from the 200 joint draws of the Stage-4a/4b system; each column's plug-in is recomputed within every draw, so the standard errors carry all cross-stage covariances, and the average-effect draws include the path-mean redraw of Corollary 2, so their intervals cover the population effects.

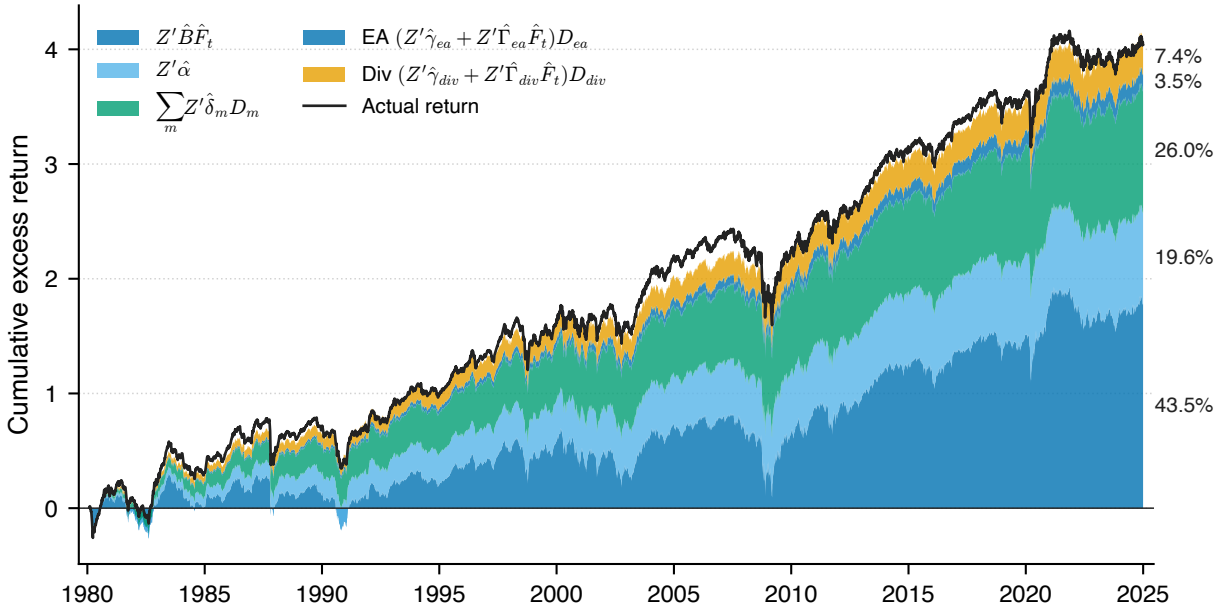


Figure 9: Cumulative equally weighted excess return and the fitted decomposition from the stacked event-type specification. The two firm-event components include both their event-level and factor-dependent terms. "Div" denotes the composite CRSP distribution ex-date event. Percentages at the right edge are shares of cumulative fitted return.

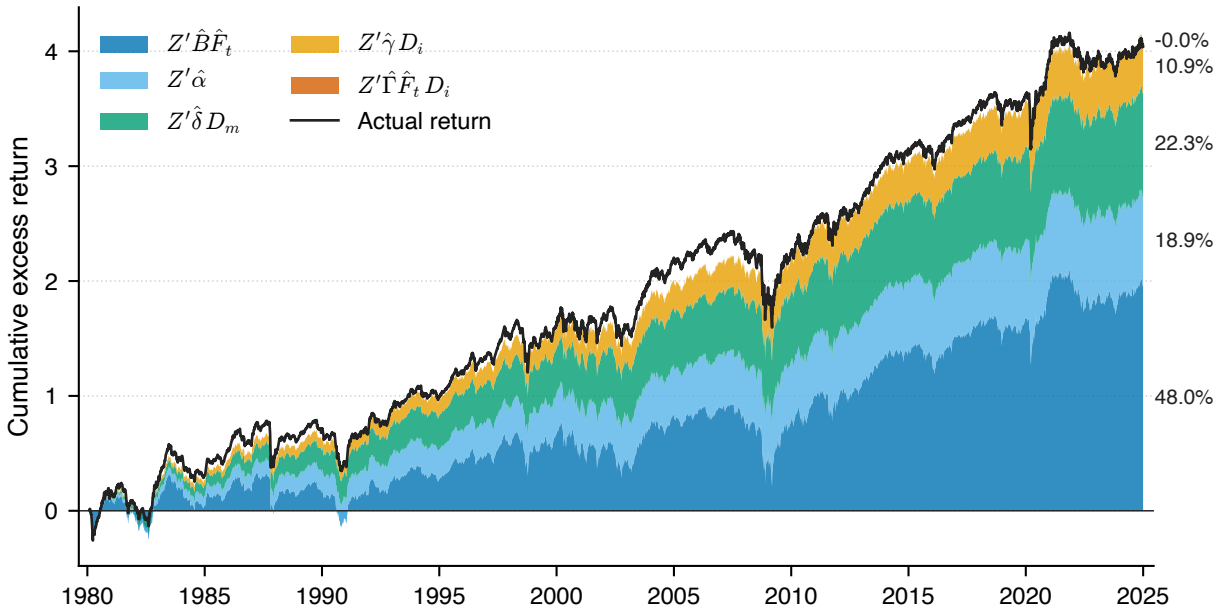


Figure 10: Cumulative equally weighted excess return and the model-implied decomposition into a baseline factor component, normalized event-free level, macro-event contribution, normalized firm-event level, and firm-event factor-dependent contribution. The CRW normalization identifies each displayed component; their sum is invariant to factor centering choices.

Figure 10 reports the decomposition for the pooled specification. The baseline factor component contributes 48.0% of cumulative fitted returns and the event-free level contributes 18.9%. Macro events contribute 22.3%, while the pooled firm-event level contributes 10.9% despite firm events accounting for only 2.2% of observations. The average firm-event factor-dependent contribution is small. The two cumulative figures use different specifications and normalize by their own cumulative fitted returns, so percentages should be compared within each figure.

4.2.3 Industry Peer-Announcement Exposure

The above described results are based on firm's own events. We now extend the analysis to understand the impact of industry peer announcements on firm returns. For such purpose, we define a bellwether as a firm in the top decile of lagged market equity within its three-digit SIC industry and month. When a bellwether announces earnings, other firms in the industry are classified as exposed. Bellwethers are firms that are more likely to have announcements that contain industry-relevant information. The announcing firm is not counted as its own peer, although it can be exposed by another bellwether in the same industry. Peer exposure occurs on 8.3% of firm-days within the data.

We re-estimate the model with earnings announcements, distribution ex-dates, and peer exposure entered jointly with the four macro-event indicators. The direct-event estimates nearly

Table 4: Average conditional-return components with peer-announcement exposure

Component	Estimate (bps)	Bootstrap SE	t-statistic
Event-free baseline	0.71^{***}	0.08	9.50
Earnings	9.55^{***}	0.80	12.00
Distribution ex-date, broad	51.95^{***}	1.18	44.05
Peer-announcement exposure (bellwether)	0.57^{**}	0.27	2.08

Notes: Characteristic-mean intercepts, in basis points per day, from the stacked specification with three firm-level event or exposure types, earnings, broad CRSP distribution ex-dates, and peer-announcement exposure, and four macro announcement types. The peer-exposure indicator equals one on the trading day after a bellwether in the same three-digit SIC industry announces earnings. Bellwethers are firms in the top decile of lagged market equity within industry and month.

unchanged when peer exposure is added. The earnings coefficient stays at 9.5 basis points, the broad-distribution coefficient at 51.9, and the FOMC coefficient moves 17 to 17.3. The stability of these estimates shows that the peer indicator captures a separate conditional-return difference.

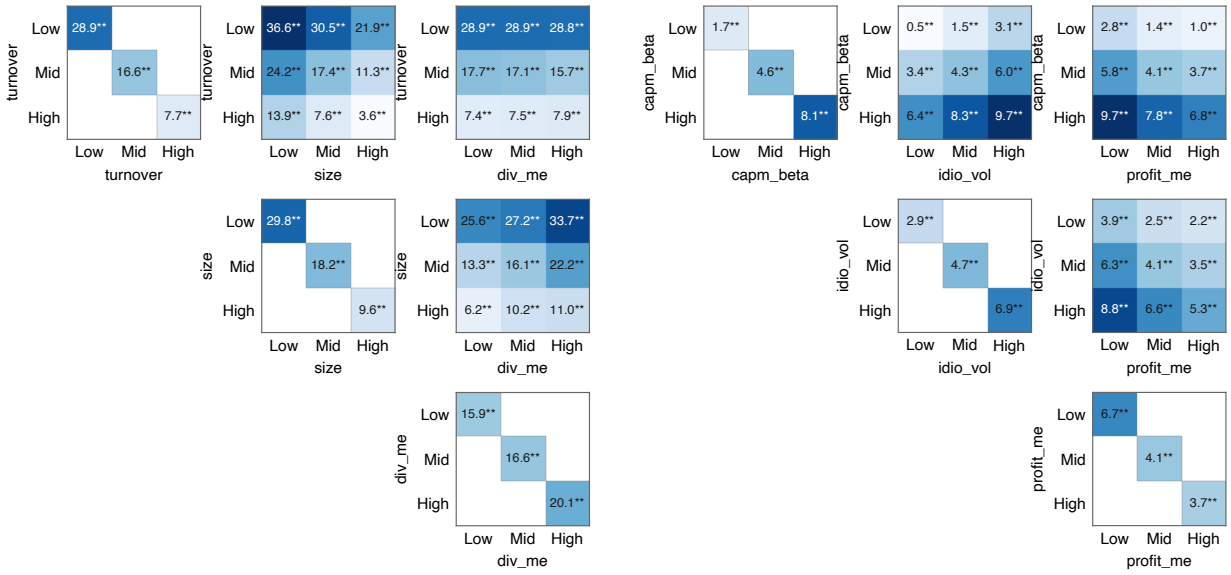
The peer-exposure coefficient is 0.57 basis points, with a baseline bootstrap standard error of 0.27. Firms in the same industry may share shocks on announcement dates, so we also use one multiplier for each three-digit SIC industry. The industry-clustered standard error is 0.33, giving a t-statistic of 1.72. The direct event estimates are largely unchanged under clustering. The characteristic slopes provide a more detailed description of the peer result. Profit-to-market, profit growth, and return on capital have the largest positive coefficients. Operating liquidity, return on equity, return on investment, and short-term reversal have negative coefficients.

Because the peer indicator records only whether a bellwether announces, the estimate should be read as an average return difference following peer news. It does not distinguish the sign or magnitude of the news, nor between one and several announcements, and observations without matched industry and size information are classified as unexposed. Accordingly, the exercise illustrates how a prespecified network exposure can be added without materially changing the direct event coefficients. It does not identify information transmission, which would require announcement surprises, exposure intensities, and industry-by-date controls (Foster, 1981; Freeman and Tse, 1992; Ramnath, 2002; Thomas and Zhang, 2008). We exclude peer exposure from the subsequent decompositions and portfolio exercises, which use the baseline specification with the two direct firm-event types.

4.3 Event Portfolios

We next describe portfolios formed within the reported event samples. Published macro calendars and many distribution ex-dates are known before the event. Compustat provides the realized earnings report date, but the data do not establish when that date became known. The earnings portfolios should therefore be read as retrospective event-conditioned portfolios.

Here the question is more granular. The broad event portfolios provide the benchmarks: a



(a) Firm events

(b) Macro events

Figure 11: Event-portfolio abnormal returns on treated observations. Entries are bucket means of the model-implied event component in basis points: $Z'\widehat{\gamma}$ for firm events and $Z'\widehat{\delta}$ for macro events. Diagonal panels are single sorts; upper triangular panels are sequential double sorts. Stars summarize the cross-date dispersion of the bucket means through the conventional time-series statistic.

portfolio of all firms with an earnings announcement or distribution ex-date and an equally weighted portfolio of all stocks on macro-announcement dates. We then ask whether the heterogeneous conditional-return patterns estimated above translate into characteristic sorts within those recorded event subsamples. For firm events, we consider turnover, size, and equity payout-to-market; for macro events, they are CAPM beta, idiosyncratic volatility, and profit-to-market. The resulting portfolios test whether the model identifies where the event premium is concentrated, beyond the profitability of simply holding the standard event portfolio.

Figure 11 first reports the model-implied event component on treated observations. The firm-event panel sorts observations by the firm-event signal $Z'\widehat{\gamma}$; the macro-event panel sorts observations by the macro-event signal $Z'\widehat{\delta}$, both including the intercept, so the pooled mean of each panel equals the corresponding event intercept by construction (18.2 and 4.8 basis points): characteristics are demeaned within date. The firm-event signal produces much larger dispersion than the macro-event signal. Across the firm-event sorts, abnormal returns are positive in all thirty-six cells and range from 3.6 to 36.6 basis points per day. The macro-event cells are also uniformly positive but much less dispersed, between 0.5 and 9.7 basis points. This contrast is consistent with the model estimates: macro announcements matter for aggregate expected returns, whereas firm events generate sharper cross-sectional variation across treated stocks.

Figure 12 evaluates the realized returns of these firm-event portfolios. The portfolios are formed

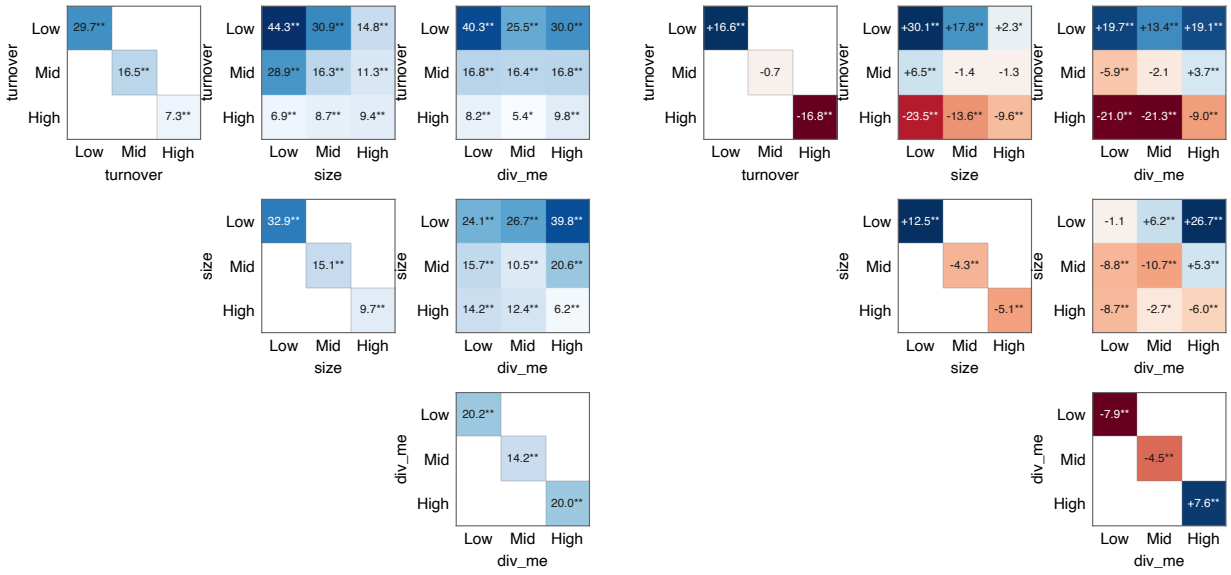
date by date using only firms with a firm event, but the characteristic bins are defined before conditioning on the event. Specifically, each trading date's full stock cross-section is split into terciles using the relevant characteristic, and firm-event stocks then inherit those tercile labels. Diagonal cells are single sorts: the bucket return is the equally weighted mean excess return of event stocks that fall in the corresponding full-sample tercile. Upper-triangular cells are sequential double sorts: outer terciles are formed on the row characteristic over the full date cross-section, inner terciles of the column characteristic are formed within each date-by-outer bucket, and event stocks inherit the two labels. The return rule is hierarchical. If a date-cell contains at least one firm-event stock, the portfolio return is the equally weighted return of those stocks. If the cell contains no firm-event stocks but the date has other firm events, the return is the standard equally weighted firm-event portfolio for that date. If the date has no firm events at all, the return is the market portfolio return. This rule keeps the portfolio time series balanced while preserving the actual event-stock return whenever a cell is populated.

The fill has a consequence for the second benchmark below. When a date-cell is empty the series takes the standard firm-event portfolio return, and that same portfolio is the regressor in Panel (b) of Figure 12; on filled dates the dependent variable therefore equals the regressor by construction, which pulls the estimated intercept toward zero and raises the regression R^2. In the rebuilt exhibits the fill share is recorded: 5.5% of date-cells use a fallback return (5.4% the standard event portfolio, 0.1% the market on the eight dates with no firm event), so the attenuation is mild but present. Panel (b) remains an attenuated comparison: it is informative about the sign pattern across cells, but its magnitudes are a lower bound on the incremental return available from sorting within the event set.

The figure then asks whether these reported portfolios earn abnormal returns relative to two benchmarks. The left panel reports alphas from regressions on the Fama–French five factors. These alphas remain large for many of the firm-event cells, showing that the event-sort returns are not absorbed by market, size, value, profitability, or investment exposures. The right panel benchmarks each portfolio against the standard firm-event portfolio. This second benchmark removes the average premium earned by event stocks on the same dates and isolates the incremental return from sorting among treated firms. Its cells fall on both sides of zero: fourteen of the thirty-six intercepts are positive, all fourteen starred, the largest at 30.1 basis points per day, and twenty-two are negative, seventeen of them starred, the most negative at -23.5. The characteristic sorts therefore locate event firms that beat the average event stock and event firms that lag it. Read together with the fill caveat above, the panel identifies where the event premium is concentrated.

4.4 Out-of-Sample Evidence

We first test whether the fitted cross-sectional relation transfers to firms that were not used to estimate it. Firms are divided into five folds. For each fold, the model is estimated on the other four and evaluated on the held-out firms. The procedure preserves the realized sequence of aggregate



(a) Alpha relative to FF5

(b) Alpha relative to standard event portfolio

Figure 12: Firm-event portfolio alphas. Entries are daily regression intercepts in basis points. Panel (a) reports alphas from regressions on the Fama–French five factors. Panel (b) reports alphas from regressions on the standard equally weighted firm-event portfolio. Diagonal panels are single sorts; upper triangular panels are sequential double sorts. Stars denote pointwise 10% and 5% significance of the intercept under Newey–West (five-lag) standard errors from the daily time-series regressions, with no multiplicity adjustment across the 36 cells; with that many cells, three to four 10% markers are expected by chance, and cell-level extremes are read descriptively.

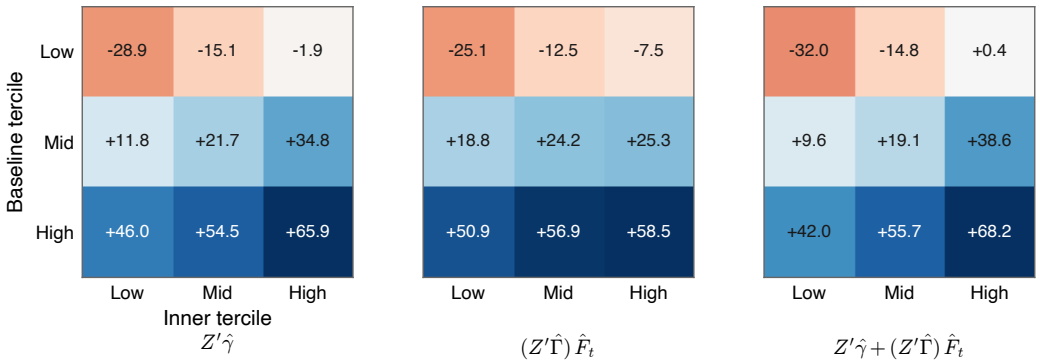


Figure 13: Out-of-sample double-sort portfolios for firm-event observations. Rows are baseline terciles, implied by $Z'(\hat{\alpha}+\hat{B}\widehat{F}_t)$. Columns are inner terciles formed using, from left to right, $Z'\widehat{\gamma}$, $Z'\widehat{\Gamma}\widehat{F}_t$, and their sum. Entries are held-out mean realized excess returns in basis points per day. The high-minus-low spreads quoted in the text are the inner-high minus inner-low differences of the date-level portfolio means underlying these cells.

shocks while excluding held-out returns from the coefficient estimates. The held-out signals use $\widehat{F}_t$, which is estimated from contemporaneous returns of the training firms. Rank breakpoints and imputation medians are also computed from the full cross-section. The exercise therefore asks whether the fitted mapping transfers to unseen firms after the date-t factor realization is estimated.

We use a double sort to determine whether the event signals contain information beyond the ordinary fitted return. Firms are first sorted within date on $Z'_{i,t-1}(\widehat{\alpha}+\widehat{B}\widehat{F}_t)$. Within each baseline tercile, event observations are sorted on the relevant increment. For firm events, the increments are $Z'\widehat{\gamma}$, $Z'\widehat{\Gamma}\widehat{F}_t$, and their sum. For macro events, the increment is $Z'\widehat{\delta}$.

Figure 13 shows a monotone relation between the firm-event scores and held-out returns. Sorting on $Z'\widehat{\gamma}$ produces high-minus-low spreads of 27.1, 23.2, and 19.9 basis points across the three baseline rows. The factor-dependent component produces smaller positive spreads of 17.6, 6.4, and 7.6 basis points. The total signal produces the widest spreads at 32.4, 28.9, and 26.1 basis points. Both components therefore contribute to the ranking.

Figure 14 repeats the sort on macro-event dates. Returns increase across the baseline rows. Within those rows, the macro signal produces high-minus-low spreads of -7.4, 4, and 15.9 basis points. Its incremental relation with returns is therefore concentrated in the middle and high baseline groups.

4.5 Out-of-Sample Event Portfolios

The previous portfolio exercises use signals estimated on the full sample, and the cross-sectional validation holds out firms but still uses the realized history of all dates when estimating the model. The final exercise imposes a stricter trading discipline: at the start of each year, the investor observes only the past, estimates the firm-event signal, and then uses that fixed signal to sort event

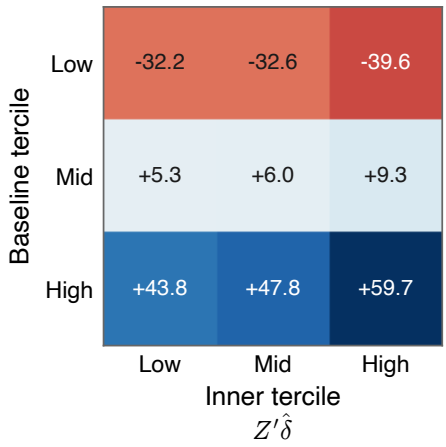


Figure 14: Out-of-sample double-sort portfolios for macro-event dates. Rows are baseline terciles, implied by $Z'(\hat{\alpha} + \hat{B}\widehat{F}_t)$. Columns are inner terciles formed using the macro-event signal $Z'\widehat{\delta}$. Entries are held-out mean realized excess returns in basis points per day.

stocks during the year. This asks whether the heterogeneous firm-event premium is not only visible in-sample, but also usable in a portfolio rule whose parameter estimation is forward-looking. We implement this test with an expanding-window design. For each target year Y, all parameters are estimated using only observations from 1980 through $Y-1$, and the resulting estimates are then held fixed for portfolio formation during year Y. The 1990 portfolios therefore use parameters estimated from 1980–1989, while the 2024 portfolios use parameters estimated from 1980–2023. This design eliminates look-ahead in the estimated coefficient $\widehat{\gamma}_Y$, but it does not eliminate look-ahead in event membership: the current implementation selects earnings observations using realized Compustat report dates.

The expanding-window estimation repeats the full four-stage procedure of Section 3.2 separately for each target year, with $K = 3$, on the stacked specification that separates earnings announcements from distribution ex-dates. For each trading date t from 1990 through 2024, every firm with a firm event is scored using the parameter vector estimated before the start of that year, taking the event-type coefficient that matches the firm's event:

$$s_{i,t} = \begin{cases} Z'_{i,t-1}\widehat{\gamma}_{\mathrm{ea},Y(t)} & \text{if firm } i \text{ announces earnings on date } t, \\ Z'_{i,t-1}\widehat{\gamma}_{\mathrm{div},Y(t)} & \text{otherwise.} \end{cases}$$

The standard event portfolio is the equally weighted excess return of all firms with a firm event on date t. When at least three firms have an event, the event firms are sorted by $s_{i,t}$ and split into low, middle, and high tercile portfolios. The high-minus-low portfolio is the difference between the high- and low-score event portfolios. When fewer than three firms have an event, all three sorted portfolios are set equal to the standard event portfolio and the high-minus-low return is set to zero. Dates with no firm-event observations are filled with market excess returns.

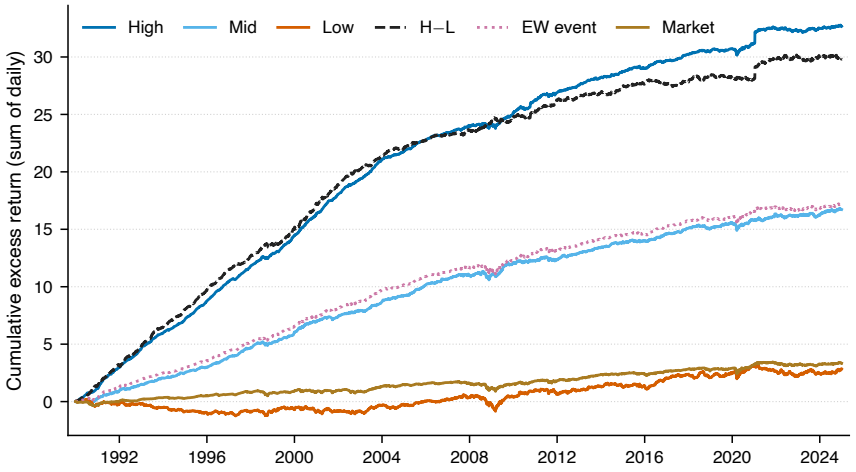


Figure 15: Cumulative daily excess returns of the expanding-window event portfolios. Each year uses firm-event parameters estimated only from prior years. The low, middle, and high portfolios sort firm-event observations by $Z'\widehat{\gamma}$; the standard event portfolio is the equally weighted return of all firm-event observations; H–L is the high-minus-low portfolio. Series are cumulative simple sums of daily excess returns; on dates with no firm-event observation the level portfolios carry the market excess return and H–L is zero.

Figure 15 plots cumulative daily excess returns. The first pattern is the standard event effect: the equally weighted portfolio of firm-event stocks outperforms the market portfolio, consistent with recurring firm events being high-return dates. The second pattern is the incremental role of characteristics. Within the set of event stocks, the portfolio with the highest out-of-sample firm-event score persistently outperforms the middle- and low-score portfolios. The high-minus-low portfolio therefore removes the common event-day premium and isolates the return from ranking treated firms using only parameters estimated from prior data. Its positive and persistent cumulative return shows that coefficients estimated from prior years continue to rank returns conditional on realized event membership.

5 Conclusion

Scheduled events change both the level and the cross-section of expected returns. This paper develops a characteristic-based latent-factor model that separates firm-event premia, macro-event premia, and common return variation. The framework accommodates several event types, overlapping event calendars, and differences in event exposure across firms.

The estimator proceeds in four stages. It first estimates the characteristic-managed return process outside firm events, uses calendar variation to estimate macro-event premia, recovers the latent factor space, and then estimates firm-event levels and factor exposures. The stages are estimated sequentially but analyzed jointly because each stage uses quantities generated by the preceding stages. The split multiplier bootstrap re-estimates the complete procedure and accounts

for the resulting estimation error. A half-panel jackknife corrects the finite-sample attenuation in the exposure estimates when the time series is long relative to the cross-section.

The empirical results show substantial differences across event types. Earnings announcements, cash-distribution ex-dates, and FOMC decisions carry positive average returns, but their characteristic slopes and factor exposures differ. The cash-distribution estimate is consistent with the classical incomplete ex-dividend price adjustment. Returns also rise modestly after bellwether firms in the same industry announce earnings, although this evidence is suggestive and does not identify the transmission of information across firms.

Event scores estimated using prior data continue to rank returns on subsequent event dates. This evidence indicates that the estimated heterogeneity persists out of sample, although the resulting daily trading strategy does not survive transaction costs. More broadly, recurring events provide useful variation for studying when expected returns arise and which firms earn them.

References

Ai, H. and Bansal, R. (2018). Risk preferences and the macroeconomic announcement premium. *Econometrica* 86, 1383–1430.

Ahn, S. C. and Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica* 81, 1203–1227.

Ali, U. and Hirshleifer, D. (2020). Shared analyst coverage: Unifying momentum spillover effects. *Journal of Financial Economics* 136, 649–675.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica* 70, 191–221.

Bai, J. and Ng, S. (2006). Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica* 74, 1133–1150.

Bai, J. and Ng, S. (2023). Approximate factor models with weaker loadings. *Journal of Econometrics* 235, 1893–1916.

Ball, R. and Kothari, S. P. (1991). Security returns around earnings announcements. *The Accounting Review* 66, 718–738.

Barber, B. M., De George, E. T., Lehavy, R., and Trueman, B. (2013). The earnings announcement premium around the globe. *Journal of Financial Economics* 108, 118–138.

Beaver, W. H. (1968). The information content of annual earnings announcements. *Journal of Accounting Research* 6, 67–92.

Berg, T., Reisinger, M., and Streitz, D. (2021). Spillover effects in empirical corporate finance. *Journal of Financial Economics* 142, 1109–1127.

Binder, J. J. (1985). On the use of the multivariate regression model in event studies. *Journal of Accounting Research* 23, 370–383.

Borusyak, K., Hull, P., and Jaravel, X. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies* 91, 3253–3285.

Brown, S. J. and Warner, J. B. (1985). Using daily stock returns: The case of event studies. *Journal of Financial Economics* 14, 3–31.

Chen, Q., Roussanov, N., and Wang, X. (2025). *Semiparametric Conditional Factor Models in Asset Pricing*. arXiv:2112.07121v5.

Chernozhukov, V., et al. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* 21, C1–C68.

Cohen, L. and Frazzini, A. (2008). Economic links and predictable returns. *Journal of Finance* 63, 1977–2011.

Daniel, K., Grinblatt, M., Titman, S., and Wermers, R. (1997). Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance* 52, 1035–1058.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110, 2964–2996.

deHaan, E., Shevlin, T., and Thornock, J. (2015). Market (in)attention and the strategic scheduling and timing of earnings announcements. *Journal of Accounting and Economics* 60, 36–55.

Eades, K. M., Hess, P. J., and Kim, E. H. (1984). On interpreting security returns during the ex-dividend period. *Journal of Financial Economics* 13, 3–34.

Elton, E. J. and Gruber, M. J. (1970). Marginal stockholder tax rates and the clientele effect. *Review of Economics and Statistics* 52, 68–74.

Engelberg, J., McLean, R. D., and Pontiff, J. (2018). Anomalies and news. *Journal of Finance* 73, 1971–2001.

Fama, E. F., Fisher, L., Jensen, M. C., and Roll, R. (1969). The adjustment of stock prices to new information. *International Economic Review* 10, 1–21.

Fama, E. F. (1998). Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics* 49, 283–306.

Fan, J., Liao, Y., and Wang, W. (2016). Projected principal component analysis in factor models. *Annals of Statistics* 44, 219–254.

Fan, J., Ke, Y., Liao, Y., and Neykov, M. (2022). Augmented factor models with applications to validating market risk factors and forecasting bond risk premia. *Journal of Econometrics* 229, 22–52.

Foster, G. (1981). Intra-industry information transfers associated with earnings releases. *Journal of Accounting and Economics* 3, 201–232.

Frank, M. and Jagannathan, R. (1998). Why do stock prices drop by less than the value of the dividend? Evidence from a country without taxes. *Journal of Financial Economics* 47, 161–188.

Frazzini, A. and Lamont, O. A. (2007). The earnings announcement premium and trading volume. NBER Working Paper 13090.

Freeman, R. and Tse, S. (1992). An earnings prediction approach to examining intercompany information transfers. *Journal of Accounting and Economics* 15, 509–523.

Gagliardini, P., Ossola, E., and Scaillet, O. (2016). Time-varying risk premium in large cross-sectional equity data sets. *Econometrica* 84, 985–1046.

Giglio, S. and Xiu, D. (2021). Asset pricing with omitted factors. *Journal of Political Economy* 129, 1947–1990.

Giglio, S., Liao, Y., and Xiu, D. (2021). Thousands of alpha tests. *Review of Financial Studies* 34, 3456–3496.

Goldsmith-Pinkham, P. and Lyu, T. (2025). Causal inference in financial event studies. Working paper, Yale University. arXiv:2511.15123.

Graham, J. R., Michaely, R., and Roberts, M. R. (2003). Do price discreteness and transactions costs affect stock returns? Comparing ex-dividend pricing before and after decimalization. *Journal of Finance* 58, 2611–2636.

Gonçalves, S. and Perron, B. (2014). Bootstrapping factor-augmented regression models. *Journal of Econometrics* 182, 156–173.

Gu, S., Kelly, B., and Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics* 222, 429–450.

Hartzmark, S. M. and Solomon, D. H. (2013). The dividend month premium. *Journal of Financial Economics* 109, 640–660.

Hartzmark, S. and Solomon, D. (2018). Recurring firm events and predictable returns: The within-firm time series. *Annual Review of Financial Economics* 10, 499–517.

Hirshleifer, D., Lim, S. S., and Teoh, S. H. (2009). Driven to distraction: Extraneous events and underreaction to earnings news. *Journal of Finance* 64, 2289–2325.

Jaffe, J. F. (1974). Special information and insider trading. *Journal of Business* 47, 410–428.

Johnson, T. L. and So, E. C. (2018). Time will tell: Information in the timing of scheduled earnings news. *Journal of Financial and Quantitative Analysis* 53, 2431–2464.

Kalay, A. (1982). The ex-dividend day behavior of stock prices: A re-examination of the clientele effect. *Journal of Finance* 37, 1059–1070.

Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134, 501–524.

Kim, S., Korajczyk, R. A., and Neuhierl, A. (2020). Arbitrage portfolios. *Review of Financial Studies* 34, 2813–2856.

Kothari, S. P. and Warner, J. B. (2007). Econometrics of event studies. In *Handbook of Corporate Finance: Empirical Corporate Finance*, vol. 1, 3–36. Elsevier.

Lo, A. W. (2002). The statistics of Sharpe ratios. *Financial Analysts Journal* 58(4), 36–52.

Liao, Y., Tong, X., Wang, W., and Xiu, D. (2026). *Fixed-Order PCA: Theory for Overestimated Factor Models*. arXiv:2605.18448.

Lucca, D. and Moench, E. (2015). The pre-FOMC announcement drift. *Journal of Finance* 70, 329–371.

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature* 35, 13–39.

Mammen, E. (1993). Bootstrap and wild bootstrap for high-dimensional linear models. *Annals of Statistics* 21, 255–285.

Mandelker, G. (1974). Risk and return: The case of merging firms. *Journal of Financial Economics* 1, 303–335.

Murphy, K. M. and Topel, R. H. (1985). Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics* 3, 370–379.

Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *Review of Economics and Statistics* 92, 1004–1016.

Onatski, A. (2012). Asymptotics of the principal components estimator of large factor models with weakly influential factors. *Journal of Econometrics* 168, 244–258.

Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International Economic Review* 25, 221–247.

Patton, A. J. and Verardo, M. (2012). Does beta move with news? Firm-specific information flows and learning about profitability. *Review of Financial Studies* 25, 2789–2839.

Ramnath, S. (2002). Investor and analyst reactions to earnings announcements of related firms: An empirical analysis. *Journal of Accounting Research* 40, 1351–1376.

Roth, J., Sant'Anna, P. H. C., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics* 235, 2218–2244.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66, 688–701.

Rubin, D. B. (1980). Comment: Randomization analysis of experimental data: The Fisher randomization test. *Journal of the American Statistical Association* 75, 591–593.

Savor, P. and Wilson, M. (2014). Asset pricing: A tale of two days. *Journal of Financial Economics* 113, 171–201.

Savor, P. and Wilson, M. (2016). Earnings announcements and systematic risk. *The Journal of Finance* 71, 83–138.

Schipper, K. and Thompson, R. (1983). The impact of merger-related regulations on the shareholders of acquiring firms. *Journal of Accounting Research* 21, 184–221.

Shanken, J. (1992). On the estimation of beta-pricing models. *Review of Financial Studies* 5, 1–33.

Thomas, J. and Zhang, F. (2008). Overreaction to intra-industry information transfers? *Journal of Accounting Research* 46, 909–940.

Thompson, R. (1985). Conditioning the return-generating process on firm-specific events: A discussion of event study methods. *Journal of Financial and Quantitative Analysis* 20, 151–168.

Internet Appendix

Event Alpha and Event Risk Premia

This Internet Appendix contains: Appendix A, the Monte Carlo simulation; Appendix B, joint limit theory and the bootstrap; Appendix C, compact proofs; and Appendix D, firm characteristics.

A Monte Carlo Simulation

The simulation studies the finite-sample behavior of the four-stage estimator. It examines the level and macro-premium estimates, the attenuation of the exposure estimates, the split firm/date bootstrap, and the half-panel jackknife of Proposition 2. The design includes two firm-event types and two macro-event types, allowing the event indicators to overlap.

A.1 Design

We simulate (14) with $J = 2$ firm-event types and $M = 2$ macro types at a fixed parameter vector. The event indicators are independent Bernoulli draws. Stage 4a uses firm-event observations on dates without a macro event, while firm–macro coincidences enter Stage 4b. The simulation reports the level, exposure, loading, and macro-premium blocks, together with the average firm-event effects.

The baseline has $L = 4$ characteristics and $K = 2$ factors. Panel C of Table 5 uses the empirical dimensions $L = 30$ and $K = 3$. Characteristics follow an AR(1) with persistence 0.97 and are rank-normalized within date. The columns of B are orthonormal, and α is projected onto their orthogonal complement. We set $\sigma_\alpha = 0.01$, $\sigma_\Gamma = 0.30$, and $U_{it} \sim N(0, 0.06^2)$. The firm-event levels have standard deviation 0.005, and each coordinate of λ_g has standard deviation equal to one-half of the corresponding factor standard deviation. We impose $\delta_g = B\lambda_g$. The two factor standard deviations are 0.04 and 0.03, with the same geometric decay when $K = 3$. All coefficients are drawn once and held fixed across replications.

Each firm-event type occurs with probability 7%, and each macro type with probability 10%. The resulting union firm-event rate of 13.5% is higher than its empirical counterpart and provides enough observations for componentwise coverage calculations. One design lowers the per-type firm-event rate to 2%. Panels A and B use $1,000$ replications and 200 bootstrap draws per replication. Panel C uses 300 replications at the higher dimension. Panel A considers $(N, T) \in \{(200, 240), (200, 120), (400, 120), (800, 120)\}$, so that T/N falls from 1.20 to 0.15. Panel B uses $(100, 200)$, $(100, 300)$, and $(200, 600)$, then repeats the last design at the lower event rate. Panel C uses $(200, 600)$ and $(800, 600)$ with $L = 30$ and $K = 3$. Its Stage-4a regressor dimension is therefore $L(K + 1) = 120$, as in the application. Each replication reruns the estimator and the bootstrap

Table 5: Finite-sample performance of the four-stage estimator

(N,T)	T/N	L	K	Event rate	Level cov.	λ z-var.	Max. Γ bias	Γ cov.	Γ cov. jk
Panel A: Declining aspect ratio									
$(200, 240)$	1.20	4	2	7%	0.93–1.00	1.18–1.29	0.044	0.78	0.91
$(200, 120)$	0.60	4	2	7%	0.93–0.99	1.33–1.43	0.044	0.85	0.91
$(400, 120)$	0.30	4	2	7%	0.94–1.00	1.28–1.45	0.021	0.90	0.92
$(800, 120)$	0.15	4	2	7%	0.93–1.00	1.25–1.40	0.011	0.91	0.93
Panel B: Full-sample aspect ratios									
$(100, 200)$	2.00	4	2	7%	0.93–0.98	1.18–1.36	0.075	0.72	0.89
$(100, 300)$	3.00	4	2	7%	0.94–0.98	1.15–1.26	0.079	0.59	0.89
$(200, 600)$	3.00	4	2	7%	0.94–0.99	1.05–1.25	0.042	0.54	0.91
$(200, 600)$	3.00	4	2	2%	0.94–0.97	1.04–1.23	0.040	0.85	0.92
Panel C: Empirical dimension, $L(K+1) = 120$									
$(200, 600)$	3.00	30	3	7%	0.89–1.00	1.11–1.36	0.043	0.57	0.83
$(800, 600)$	0.75	30	3	7%	0.92–1.00	0.97–1.19	0.011	0.79	0.78

Notes: Panels A and B use 1,000 replications and Panel C uses 300. Each replication uses 200 bootstrap draws. The Monte Carlo standard error of a coverage estimate at 0.95 is 0.007 in Panels A and B and 0.013 in Panel C. Reported ranges are taken across characteristic components, factor coordinates, and event types. Event rate is the probability of each firm-event type. The two $(200, 600)$ rows in Panel B differ only in this probability. Level cov. reports the range of componentwise coverage for nominal 95% intervals for $\widehat{\alpha}$, $\widehat{\gamma}_j$, and $\widehat{\gamma}_a$. The upper end of this range reflects conservative intervals for the announcement-day levels. The column labeled λ z-var. reports the variance of the bootstrap-standardized $\widehat{\boldsymbol{\lambda}}_g$, with target one. Max. Γ bias is the largest average absolute bias among the columns of $\widehat{\Gamma}_j$ and $\widehat{\Gamma}_a$, where the average is taken across the L characteristics. The last two columns report the lowest componentwise coverage among the same exposure estimates. The final column uses the $S = 8$ half-panel jackknife of Proposition 2. Its intervals are centered at the corrected estimates and use the full-sample bootstrap standard errors.

of Section 3.3. The bootstrap uses exponential firm weights, a Rademacher date multiplier, the firm-only Stage-2 draw in Stage 3, and the draw-specific rotation $H^\star$.

A.2 Results

Table 5 reports three sets of results. The first is coverage for the level coefficients. The second is the variance of the bootstrap-standardized structural macro-premium estimates, for which the target is one. The remaining columns report exposure bias and the worst componentwise coverage before and after the half-panel jackknife.

Coverage for the level coefficients is generally close to its nominal value. In the low-dimensional designs, the lower endpoint is between 0.93 and 0.94. Some announcement-day intervals are conservative, which accounts for upper endpoints near one. At the empirical dimension, the lower endpoint rises from 0.89 to 0.92 when N increases from 200 to 800.

The simulations also reproduce the different convergence rates of the parameter blocks. Holding $N = 200$, an increase in T from 120 to 600 reduces the RMSE of $\widehat{\boldsymbol{\lambda}}_g$ by a factor of about 2.3. This is close to the $\sqrt{5}$ reduction implied by its $\sqrt{T}$ rate. Increasing N with T fixed has little effect on this block. By comparison, the RMSE of $\widehat{\gamma}_j$ falls by about $\sqrt{2}$ when either dimension doubles.

The exposure estimates are more sensitive to the size of the cross-section. In Panel A, the

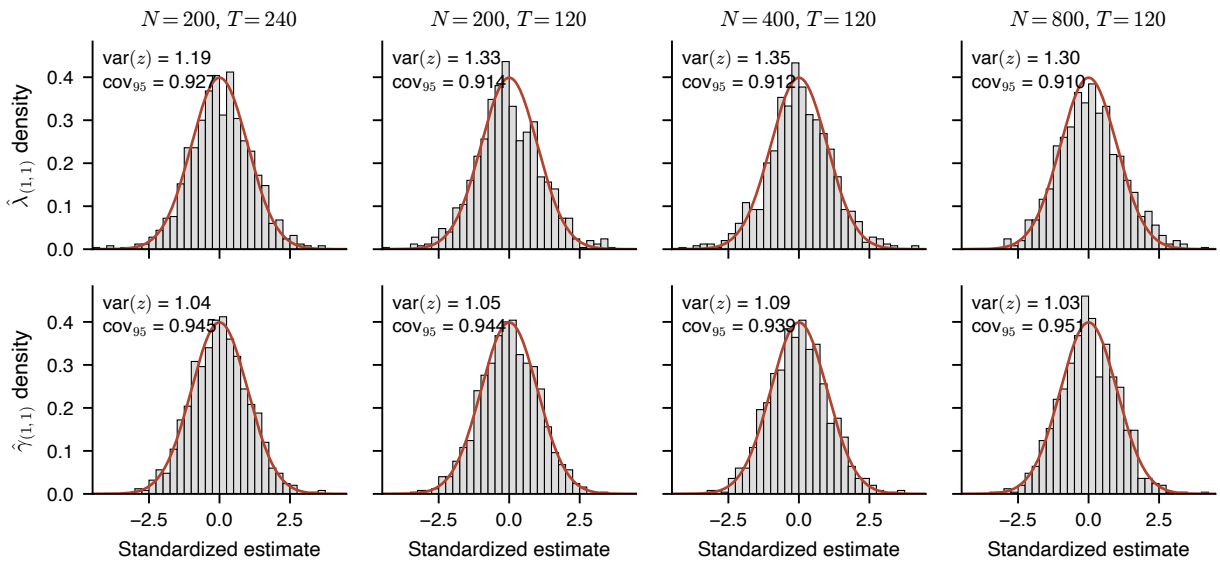


Figure 16: Standardized Monte Carlo distributions. The top panel reports one coordinate of the first structural premium vector for the sample-size configurations in Panel A of Table 5. The bottom panel reports one component of the first firm-event coefficient. The red curve is the standard normal density.

largest average bias falls from 0.044 to 0.011 as T/N falls from 1.20 to 0.15. At the larger aspect ratios in Panel B, the lowest coverage falls to 0.54. The half-panel jackknife raises this coverage to between 0.89 and 0.93 in the low-dimensional designs. Its performance is less uniform at the empirical dimension. Coverage rises from 0.57 to 0.83 when $(N, T) = (200, 600)$, but changes from 0.79 to 0.78 when $(N, T) = (800, 600)$. The correction is therefore most useful when exposure bias is substantial.

The bootstrap standard errors for the structural macro premia are somewhat too small. The largest z-variance is 1.45, implying an understatement of about 17% relative to the Monte Carlo standard deviation. In Panel C, the largest z-variance is 1.36 when $N = 200$ and 1.19 when $N = 800$. These values imply understatements of about 14% and 8%, respectively.

Separate calculations examine the average event effects. With the factor-path means redrawn, coverage ranges from 0.91 to 0.98 on dates without macro announcements and from 0.92 to 0.98 on announcement dates. If the path-mean draw is omitted, announcement-day coverage falls to 0.40 in the largest high-dimensional design.

Figure 16 reports representative standardized distributions from Panel A. The firm-event statistic is close to the standard normal benchmark. The structural-premium statistic is more dispersed, consistent with the z-variances in the table. Panel C gives a similar comparison at the empirical dimension. Increasing N from 200 to 800 raises the lower endpoint of level coverage from 0.89 to 0.92, lowers the largest z-variance from 1.36 to 1.19, and reduces the largest exposure bias from 0.043 to 0.011.

Table 6: Overestimating the number of factors

(N, T)	L	K_0	K	$\|\mathbb{E}\widehat{\alpha}\| / \|\alpha\|$	$\widehat{\alpha}$ cov.	$\widehat{\gamma}$ cov.	$\widehat{\delta}$ RMSE
Panel A: Empirical dimension							
$(800, 600)$	30	3	3	1.000	0.92–0.96	0.93–0.98	1.472×10^{-3}
$(800, 600)$	30	3	4	0.959	0.91–0.97	0.93–0.97	1.472×10^{-3}
$(800, 600)$	30	3	5	0.927	0.91–0.97	0.90–0.98	1.472×10^{-3}
Panel B: Small L							
$(200, 600)$	4	2	2	1.001	0.95	0.94–0.97	3.642×10^{-3}
$(200, 600)$	4	2	3	0.509	0.60–0.84	0.73–0.93	3.642×10^{-3}
$(200, 600)$	4	2	4	0.000	0.00	0.05–0.89	3.642×10^{-3}

Notes: K_0 is the number of factors in the simulated data, and K is the dimension supplied to Stage 3. Rows with $K = K_0$ are correctly specified. Panel A uses 300 replications and Panel B uses 1,000. Each replication uses 200 bootstrap draws. Coverage entries are componentwise ranges for nominal 95% intervals. The RMSE of $\widehat{\delta}$ is pooled across components and macro-event types.

A.3 Factor-Dimension Misspecification

The final experiment varies the number of factors supplied to the estimator. The data contain K_0 factors, while Stage 3 is estimated with $K \geq K_0$. The additional directions contain no time-varying factor risk. Table 6 reports results for the empirical dimension and for a design with $L = 4$.

The RMSE of $\widehat{\delta}$ is identical across values of K because Stage 2 is completed before the factor dimension is imposed. The effect of increasing K instead appears in $\widehat{\alpha}$ and the firm-event coefficients. At the empirical dimension, increasing K from three to five reduces $\|\mathbb{E}\widehat{\alpha}\| / \|\alpha\|$ from 1.000 to 0.927, while coverage for both level blocks remains near its nominal value.

The consequences are much larger when $L = 4$. Estimating three instead of two factors reduces $\|\mathbb{E}\widehat{\alpha}\| / \|\alpha\|$ to 0.509, and coverage deteriorates for both $\widehat{\alpha}$ and $\widehat{\gamma}$. When $K = L = 4$, the estimated loading space spans the full characteristic space and forces $\widehat{\alpha}$ to zero. The exercise indicates that moderate overextraction is less consequential when the characteristic dimension is large, but it can be severe when the estimated loading space nearly exhausts that dimension.

B Main Theoretical Results

B.1 Identification

We first state the identification conditions from Section 3.1 in formal terms. Proposition 1 and all subsequent results refer to these conditions by the labels we introduce here. Define

$$\mathcal{G}_{t-1} = \sigma(\mathcal{I}_{t-1}, \bar{F}_1, \ldots, \bar{F}_{t-1}).$$

This filtration contains more information than $\mathcal{I}_{t-1}$. At finite N, the latent factor $\bar{F}_{t-1}$ is recovered from the date-$(t-1)$ cross-section with estimation error and is therefore not measurable with respect to $\mathcal{I}_{t-1}$. Conditioning on $\mathcal{G}_{t-1}$ makes the date-t factor and calendar scores martingale differences

relative to a filtration that contains their past values. This property allows for averaging the calendar score over dates.

Assumption 1 (Sampling, timing, and rank)**.**

(i) Event indicators and characteristics are $\mathcal{I}_{t-1}$-measurable,

$$\mathbb{E}[U_{it}(d) \mid \mathcal{G}_{t-1}, \bar{F}_t] = 0, \qquad \mathbb{E}[\bar{F}_t \mid \mathcal{G}_{t-1}] = \mu, \quad d \in \{0, 1\},$$

and $\mathbb{E}\|\bar{F}_t\|^2 < \infty$ for every t.

(ii) The control set is nonempty on every date, and with $N_t^{(0)} := |\{i : D_{i,t-1} = 0\}|$, the realized control Gram matrix $N_t^{(0)-1} \sum_{i:D_{i,t-1}=0} Z_{i,t-1} Z_{i,t-1}'$ is nonsingular for every t, almost surely.

(iii) With $X_{it}^0 := (Z_{i,t-1}', (Z_{i,t-1} \otimes \bar{F}_t)')'$, the clean-event second moment $\mathbb{E}[X_{it}^0 X_{it}^{0\prime} \mid D_{i,t-1} = 1,\ D_{m,t-1} = 0]$ exists and is nonsingular, and for some t, $\mathbb{P}(D_{i,t-1} = 1, D_{m,t-1} = 0) > 0$. The calendar is nonconstant: $\mathbb{P}(D_{m,t-1} = 1) > 0$ for some t and $\mathbb{P}(D_{m,t-1} = 0) > 0$ for some t.

With multiple event types, condition (iii) applies to the stacked objects. The variable X_{it}^0 is replaced by the indicator-interacted design, which must be nonsingular. Calendar nonconstancy requires joint nondegeneracy. There must be no $c \in \mathbb{R}^M \setminus \{0\}$ and constant c_0 such that $c' \boldsymbol{D}_{t-1}^m = c_0$ almost surely for every t. When $M = 1$, this condition reduces to the stated nonconstancy condition.

Assumption 2 (Factor structure and normalization)**.**

(i) Factor structure. *The number of factors K is fixed and does not grow with N or T, and the loading matrix has rank K. The population centered factor second moment*

$$\Sigma_F := \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}[(\bar{F}_t - \mu)(\bar{F}_t - \mu)']$$

exists and is positive definite: a limit of population moments, which may vary with t.

(ii) Orthogonality normalization. $\alpha' B = 0$ *(Chen et al., 2025).*

(iii) Factor coordinates. $B'B = I_K$, *and Σ_F is diagonal with strictly decreasing diagonal entries.*

Identification is defined relative to the joint law of the panel and the latent factor path, which we call the panel law. A parameter is identified if any two parameter and factor path configurations that generate the same conditional moments given $(\mathcal{G}_{t-1}, \bar{F}_t)$ are identical up to the stated column signs. The conditional law of large numbers in Assumption 3(i) links these conditional moments to observable cross-sectional averages for estimation. Assumptions 1 and 2(i) impose only the conditions needed for identification.

Proposition 1 (Identification). *Suppose Assumptions 1 and 2(i) hold.*

(i) The panel law identifies the macro premium δ, the loading space through its projector $P_B := B(B'B)^{-1}B'$, the average slope $\bar{a} = \alpha + B\mu$, the average firm-event effect $\gamma + \Gamma\mu$, the centered factor components $B(\bar{F}_t - \mu)$ and $\Gamma(\bar{F}_t - \mu)$, and every fitted conditional return outside the coincidence cell. It does not identify α, γ and μ separately.

(ii) If Assumption 2(ii) also holds, then $\alpha = (I_L - P_B)\bar{a}$ and $B\mu = P_B\bar{a}$ are identified, and with them γ, $\Gamma\mu$, $B\bar{F}_t$ and $\Gamma\bar{F}_t$. None of these depends on the factor coordinate system.

(iii) If Assumption 2(iii) also holds, then B, Γ and $\{\bar{F}_t\}$ are identified individually, up to a common sign in each factor column.

(iv) If $\mathbb{P}(D_{i,t-1} = D_{m,t-1} = 1) > 0$ for some t and the coincidence-cell second moment $\mathbb{E}[X^0_{it}X^{0\prime}_{it} \mid D_{i,t-1} = D_{m,t-1} = 1]$ exists and is nonsingular, the coincidence coefficients (c, Γ_a) are identified on the same terms as (γ, Γ) in parts (i)–(iii).

(v) Under the factor-response model of Section 2.2, where (10) gives $\delta = B\lambda$, part (ii) identifies in addition the factor path $BF_t = B\bar{F}_t + \delta D_{m,t-1}$ and the structural coincidence intercept $\gamma_a = c - \Gamma_a\lambda$ on the coincidence design of part (iv), and part (iii) identifies $\lambda = B'\delta$ and $F_t = \bar{F}_t + \lambda D_{m,t-1}$, up to the same column signs.

None of parts (i) through (v) identifies a date-specific contrast in announcement status from the panel law. On an announcement date, the counterfactual factor $F_t(0)$ cannot be separated from the announcement surprise η_t in (5).

B.2 Joint Large-Sample Theory

The estimator is computed in four stages, which are presented separately in the main text. For inference, the four stages constitute a single generated-regressor problem. We therefore establish their joint properties in one theorem.

Throughout this appendix, the number of factors K is treated as known. Let $n_{1t} := \#\{i : (i,t) \in \mathcal{T}^{(1,0)}\}$ denote the number of clean firm-event observations on date t, and let $n_1 := \sum_{t=1}^{T} n_{1t}$ denote the pooled clean-event sample size. The number of control firms on date t is $N_t^{(0)} := \#\mathcal{T}_t^{(0)}$. We define the centered macro calendar as $\widetilde{d}_t := D_{m,t-1} - \bar{D}_m$ and its limiting second moment as $v_m := \operatorname{plim} T^{-1}\sum_t \widetilde{d}_t^2$. The Stage 1 managed-portfolio estimation error is $\widetilde{\epsilon}_t := \widehat{a}_t - a_t$.[6]

Assumption 3 (Design, innovations, and sampling).

(i) Design. *With $\widehat{Q}_{0t} := (N_t^{(0)})^{-1}\sum_{i:D_{i,t-1}=0} Z_{i,t-1}Z'_{i,t-1}$, there are matrices Q_{0t} with*

$$\sup_{t\le T}\left\|\widehat{Q}_{0t} - Q_{0t}\right\| \xrightarrow{p} 0, \qquad \inf_t \lambda_{\min}(Q_{0t}) > 0, \qquad \sup_t \|Q_{0t}\| < \infty,$$

[6] For each date, this error is $O_p(N^{-1/2})$. It also satisfies the average mean-square bound $T^{-1}\sum_t \|\widetilde{\epsilon}_t\|^2 = O_p(N^{-1})$, which is used below. The assumptions do not require its maximum over dates to be $O_p(N^{-1/2})$.

the control-firm score obeys the conditional law of large numbers

$$\frac{1}{N_t^{(0)}} \sum_{i:D_{i,t-1}=0} Z_{i,t-1} U_{it}(0) \xrightarrow{p} 0 \qquad \text{on every date as } N_t^{(0)} \to \infty,$$

$\min_t N_t^{(0)} \to \infty$, *the control fraction is nonvanishing,* $\liminf \min_{t \le T} N_t^{(0)}/N > 0$, *and* $\max_t N_t^{(0)} / \sum_s N_s^{(0)} \to 0$. *The clean-event Gram limit*

$$\Sigma_4 := \operatorname*{plim}_{N,T \to \infty} \frac{1}{n_1} \sum_{(i,t) \in \mathcal{T}^{(1,0)}} X_{it}^0 X_{it}^{0\prime}$$

exists and is positive definite, as does its stacked analogue when several event types are estimated jointly.[7]

(ii) Innovations and date scores. *For some $\varepsilon > 0$, the bivariate potential-innovation arrays are conditionally independent across firms given $(\mathcal{G}_{t-1}, \bar{F}_t)$ and satisfy a Lindeberg condition. Their moments satisfy*

$$\sup_{i,t} \mathbb{E}|U_{it}(d)|^{4+\varepsilon} < \infty, \qquad \sup_t \mathbb{E}\left\|\bar{F}_t\right\|^{4+\varepsilon} < \infty, \qquad d \in \{0,1\}.$$

In addition, $\sup_{i,t} \mathbb{E}[U_{it}(d)^2 \mid \mathcal{G}_{t-1}, \bar{F}_t] \le C$ *almost surely. The arrays* $\{\bar{F}_t - \mu\}_t$ *and* $\{D_{m,t-1}(\bar{F}_t - \mu)\}_t$ *are martingale difference sequences with respect to* $\mathcal{G}_{t-1}$.[8]

Stack the factor path by date as $\bar{F} := (\bar{F}_1, \ldots, \bar{F}_T)' \in \mathbb{R}^{T \times K}$, *and define* $M_T := I_T - T^{-1}\mathbf{1}_T\mathbf{1}_T'$. *The demeaned sample second moment satisfies*

$$T^{-1}\bar{F}' M_T \bar{F} = \Sigma_F + O_p(T^{-1/2}).$$[9]

(iii) Sampling shares and calendar moments. *The clean-event sample satisfies*

$$\frac{n_1}{NT} \to \pi_1 \in (0,1), \qquad \max_{t \le T} \frac{n_{1t}}{n_1} \to 0.$$

Define $\bar{D}_m := T^{-1}\sum_t D_{m,t-1}$. *The centered calendar second moment* $v_m := \operatorname{plim} T^{-1}\sum_{t \le T}(D_{m,t-1} - \bar{D}_m)^2$ *exists and is strictly positive. The calendar mean* $\bar{D}_m$ *converges in probability to* $\pi_m \in [0,1)$. *With multiple macro types,* v_m *denotes the centered* $M \times M$ *second-moment matrix of* $\boldsymbol{D}_{t-1}^m$, *which must be nonsingular.*

Each reported firm-event type has a nonvanishing share of the clean cell. For $n_1^{(j)} := \#\{(i,t) \in$

[7] Because $\widehat{X}_{it}$ is the feasible, aligned version of X_{it}^0, the corresponding feasible Gram matrix has the same nonsingularity property with probability approaching one.

[8] The demeaned scores used below are linear combinations of these arrays. The sample mean $\bar{D}_m$ enters only through its probability limit π_m from part (iii).

[9] This rate holds when the conditional variance process is constant or stationary and weakly dependent. It is the rate used to establish the rotation bounds in the proofs.

$\mathcal{T}^{(1,0)} \mid D_{i,t-1}^{(j)} = 1\}$, *we require* $n_1^{(j)}/(NT) \to \pi_1^{(j)} \in (0,1)$.[10]

Assumption 4 (Eigenspace regularity and the joint score limit)**.**

(i) Let $\mathcal{P}_K$ *carry a symmetric* $L \times L$ *matrix to its ordered top-*K *eigenspace projector. There are a neighborhood* $\mathcal{W}$ *of* $S_0 := B\Sigma_F B'$ *and a constant* $C < \infty$ *with* $\mathbb{P}(\widehat{A}^0 M_T \widehat{A}^{0\prime}/T \in \mathcal{W}) \to 1$ *such that* $\mathcal{P}_K$ *is Fréchet differentiable on* $\mathcal{W}$*, with derivative* $D\mathcal{P}_{K,S}$ *satisfying*

$$\sup_{S \in \mathcal{W}} \|D\mathcal{P}_{K,S}\| \le C, \qquad \|\mathcal{P}_K(S+\Theta) - \mathcal{P}_K(S) - D\mathcal{P}_{K,S}[\Theta]\| \le C\,\|\Theta\|^2$$

(ii) The stacked influence array generated by the calendar score, the control-firm managed-portfolio scores, the ordered eigenspace map, and the clean-event score obeys a joint Lindeberg–Feller central limit theorem with a finite covariance limit. The joint influence array includes the fixed-dimensional stacked score arrays, the $\sqrt{NT}$*-scaled orthogonal component of the calendar-noise score used by the* α *block, and its in-span counterpart used by Corollary 2.*

The uniform convergence and cross-sectional law of large numbers in Assumption 3(i) complement the almost-sure rank condition in Assumption 1(ii). Identification requires the realized control design to have full rank. Estimation additionally requires the sample design matrices to converge uniformly over dates. The martingale difference conditions in Assumption 3(ii) follow from Assumption 1(i), since the macro calendar is predetermined and $\mathbb{E}[\bar{F}_t \mid \mathcal{G}_{t-1}] = \mu$. The same part of the assumption imposes conditional independence of the potential innovations across firms. This condition is stronger than weak cross-sectional dependence. Theorem 1 therefore does not allow industry or other cross-firm clustering in the innovations.

Because L is fixed and Assumption 2(iii) imposes strictly separated eigenvalues of Σ_F, the projector properties in Assumption 4(i) follow from finite-dimensional perturbation theory. The condition specifies the neighborhood $\mathcal{W}$ and the constant C without adding an economic restriction. It does not impose an expansion for the ordered eigenvector representative paired with the regression alignment H in (22). Assumption 4(ii) is a high-level condition. It imposes a joint central limit theorem for the calendar score, the control-firm managed-portfolio scores, the ordered eigenspace map, and the clean-event score. The joint influence array also contains the fixed-dimensional stacked score arrays, the $\sqrt{NT}$-scaled orthogonal component of the calendar-noise score used by the α block, and its in-span counterpart used by Corollary 2. Stating the condition jointly preserves the covariance across the four stages of the estimator.

Deriving the limits in Assumption 4(ii) from lower-level sampling and moment conditions would require a martingale-array central limit theorem for the double-indexed score in (25). Natural sufficient conditions include the uniformly bounded $(4+\varepsilon)$ moments in Assumption 3(ii), $\min_t N_t^{(0)} \to \infty$, and the requirement that no date contribute a nonvanishing share of the control

[10] These limits determine the type-specific scaling in the stacked result of Theorem 1.

or clean-event sample. It would also require the sample second-moment matrix to enter $\mathcal{W}$ with probability approaching one and the quadratic remainder bound to hold uniformly on that neighborhood. We impose the joint limit directly because these requirements concern the full sequential estimator.

Collect the four parameter blocks and their estimates, the latter in the rotated coordinates that H of (22) aligns.

Let

$$\theta_\gamma^0 := \begin{pmatrix} \gamma \\ \mathrm{vec}(\Gamma') \end{pmatrix}, \qquad \widehat{\theta}_\gamma^H := \begin{pmatrix} \widehat{\gamma} \\ \mathrm{vec}((\widehat{\Gamma} H^{-1})') \end{pmatrix},$$

and define

$$\vartheta := \left(\delta', \alpha', \mathrm{vec}(B)', \theta_\gamma^{0\prime}\right)', \qquad \widehat{\vartheta} := \left(\widehat{\delta}', \widehat{\alpha}', \mathrm{vec}(\widehat{B} H^{-1})', \widehat{\theta}_\gamma^{H\prime}\right)',$$

where H aligns the estimated and population factor coordinates. Let

$$\mathcal{R}_{NT} := \mathrm{diag}\left(\sqrt{T} I_L, \sqrt{NT} I_L, \sqrt{NT} I_{LK}, \sqrt{n_1} I_{L(K+1)}\right).$$

Theorem 1 (Joint limit of the four-stage estimator)**.** *Under Assumptions 1, 2, and 3,*

$$\widehat{\vartheta} \xrightarrow{p} \vartheta.$$

If, in addition, Assumption 4 holds and $T/N \to 0$, then

$$\mathcal{R}_{NT}(\widehat{\vartheta} - \vartheta) \;\Rightarrow\; \mathcal{N}(0, \mathcal{V}), \tag{24}$$

where the covariance matrix $\mathcal{V}$ is defined blockwise below. In particular, the δ block has rate $\sqrt{T}$, while the remaining finite-dimensional blocks have their displayed $\sqrt{NT}$ or $\sqrt{n_1}$ rates. The matrix $\mathcal{V}$ is singular; its degenerate directions are described below. With M macro calendars and J firm-event types the stacked version replaces δ by $\mathrm{vec}(\boldsymbol{\delta}) \in \mathbb{R}^{LM}$ and θ_γ^0 by $\boldsymbol{\theta}_\gamma^0 \in \mathbb{R}^{JL(K+1)}$, with $\sqrt{T} I_{LM}$ and $\mathrm{diag}_{j \le J}\left(\sqrt{n_1^{(j)}}\, I_{L(K+1)}\right)$, with the type-$j$ clean-cell sizes $n_1^{(j)}$ and their limiting shares as in Assumption 3, replacing the corresponding blocks of $\mathcal{R}_{NT}$; the α and B blocks are unchanged, because Stages 1 and 3 do not depend on the number of types.

The covariance matrix $\mathcal{V}$ differs from the usual Stage 4a OLS covariance because the estimated benchmark and factor realization are common to all event observations on a given date. Define

$$X_{it}^H = \begin{pmatrix} Z_{i,t-1} \\ Z_{i,t-1} \otimes H^{-1} \bar{F}_t \end{pmatrix}, \qquad \Sigma_4 = \mathrm{plim}\, \frac{1}{n_1} \sum_{(i,t) \in \mathcal{T}^{(1,0)}} X_{it}^H X_{it}^{H\prime},$$

where Σ_4 is positive definite by Assumption 3(i). The proof of Theorem 1 shows that every block

of $\mathcal{R}_{NT}(\widehat{\vartheta} - \vartheta)$ is a fixed bounded linear transformation of

$$\begin{aligned} S_\delta &= \frac{1}{v_m\sqrt{T}} \sum_{t=1}^{T} \widetilde{d}_t B(\bar{F}_t - \mu), & S_\epsilon &= \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \sqrt{N}\, \widetilde{\epsilon}_t, \\ S_\delta^\perp &= \frac{1}{v_m\sqrt{T}} \sum_{t=1}^{T} \widetilde{d}_t (I_L - BB')\sqrt{N}\, \widetilde{\epsilon}_t, & S_B &= \mathcal{L}_B \left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} \sqrt{N}\, \widetilde{\epsilon}_t (\bar{F}_t - \mu)' \right), \\ S_\gamma &= \frac{1}{\sqrt{n_1}} \sum_{(i,t)\in\mathcal{T}^{(1,0)}} \psi_{it}, \end{aligned} \tag{25}$$

where $\mathcal{L}_B$ is the bounded eigenspace derivative and

$$\psi_{it} = X_{it}^H \{U_{it}(1) + r_{it}\}, \qquad r_{it} = -Z'_{i,t-1}\widetilde{\epsilon}_t - Z'_{i,t-1}\Gamma H(\widehat{F}_t - H^{-1}\bar{F}_t).$$

The two terms in r_{it} are the benchmark and factor estimation errors. For $\mathcal{D}P_B[G] = (I_L - BB')GB' + BG'(I_L - BB')$, the covariance in Theorem 1 is

$$\mathcal{V} = \lim \mathrm{Var} \begin{pmatrix} S_\delta \\ (I_L - BB')S_\epsilon - \mathcal{D}P_B[S_B](\alpha + B\mu) - \pi_m S_\delta^\perp \\ \mathrm{vec}(S_B) \\ \Sigma_4^{-1} S_\gamma \end{pmatrix}. \tag{26}$$

The four rows correspond to the δ, α, $\mathrm{vec}(B)$, and θ_γ blocks. Because the same score arrays enter several blocks, the off-diagonal covariances generally do not vanish. The Stage 4a diagonal block is

$$\mathcal{V}_\gamma = \Sigma_4^{-1} \left[\lim \mathrm{Var} \left\{ \frac{1}{\sqrt{n_1}} \sum_{(i,t)\in\mathcal{T}^{(1,0)}} \psi_{it} \right\} \right] \Sigma_4^{-1},$$

which reduces to the usual OLS covariance only when the benchmark and factors are observed.

The covariance matrix $\mathcal{V}$ is singular for two reasons. First, the $\sqrt{T}$-scaled limit of $\widehat{\delta} - \delta$ lies in $\mathrm{col}(B)$, so its covariance has rank K. The remaining $L - K$ directions are not unidentified. The span residual $(I_L - BB')\widehat{\delta}$ varies at the faster $\sqrt{NT}$ rate. Second, Stage 3 imposes $\widehat{\alpha}'\widehat{B} = 0$ exactly. If $\mathcal{G}_\alpha$ and $\mathcal{G}_B$ are the limits of $\sqrt{NT}(\widehat{\alpha} - \alpha)$ and $\sqrt{NT}(\widehat{B}H^{-1} - B)$, then

$$B'\mathcal{G}_\alpha + \mathcal{G}_B'\alpha = 0. \tag{27}$$

The stacked limit is therefore degenerate in $L-K$ directions of the δ block and in K joint directions of the (α, B) blocks. These restrictions follow from the mixed rates and the normalization. Full-vector studentization would require a generalized inverse, and none of the results below requires $\mathcal{V}^{-1}$.

Remark 1 (Rotation and reported objects). The loading and exposure estimators are compared with BH and ΓH, where H aligns the estimated and population factor coordinates. This unobserved

rotation equals a signed identity matrix up to an $O_p(T^{-1/2})$ rotation within col(B). The limit in Theorem 1 therefore has a coordinate-independent interpretation for functionals invariant to H. These include fitted conditional returns, the average firm-event effect $z'(\gamma + \Gamma\mu)$, the event alpha $z'\gamma$, the macro premium $z'\delta$, the span residual $(I_L - BB')\delta$, and functions of BB'. Individual entries of B, Γ, and $\widehat{\lambda}_g = \widehat{B}'\widehat{\delta}_g$ depend on the Stage 3 sign and rotation convention. Their standard errors are conditional on that convention. Under $\delta_g = B\lambda_g$, the coordinate-independent factor announcement premium is $\widehat{B}\widehat{\lambda}_g = \widehat{P}_B\widehat{\delta}_g$, and

$$\sqrt{T}(\widehat{B}\widehat{\lambda}_g - B\lambda_g) = P_B\sqrt{T}(\widehat{\delta}_g - \delta_g) + o_p(1).$$

Rotation invariance does not by itself establish coverage. Formal inference is governed by Corollaries 1 and 2 and by Proposition 3.

Remark 2 (Path means in average event effects). The parameter vector ϑ does not contain the factor-path mean. Let μ_T^0 denote the intercept from the time-series projection of $\bar{F}_t$ on the macro calendars. With one macro calendar, it is the average of $\bar{F}_t$ over dates without a macro announcement. The estimator

$$\widehat{\mu} = \widehat{B}'\widehat{\bar{a}}^{\,0} = \frac{1}{T}\sum_{t=1}^{T}\widehat{F}_t$$

recovers μ_T^0 in the Stage 3 coordinates up to an $O_p((NT)^{-1/2} + N^{-1})$ error generated by the scores in (25). The average event effect $\tau^\gamma(z) = z'(\gamma + \Gamma\mu)$ differs from $z'(\gamma + \Gamma\mu_T^0)$ by $z'\Gamma(\mu_T^0 - \mu)$, which is $O_p(T^{-1/2})$ under Assumption 1(i) and lies outside the influence representation of ϑ. The announcement-day effect also depends on the type-specific calendar mean of the factor path, whose order is $O_p(T_g^{-1/2})$. The bootstrap must therefore redraw these path means in addition to the components of ϑ.

B.3 Bootstrap

Each bootstrap draw repeats all four estimation stages. Mean-one firm weights are applied in Stages 1, 4a, and 4b. An independent mean-zero date multiplier perturbs the centered Stage 2 score. Stage 3 then re-estimates and aligns the loading space before Stages 4a and 4b are recomputed. The same firm weight is applied to every observation for a given firm. The bootstrap therefore preserves arbitrary within-firm dependence in the Stage 1 and Stage 4a scores. The date multiplier is applied to the calendar score. Date-level variation enters the other stages through the estimated benchmark and factor, which are common across firms on a given date and are recomputed in every draw. The bootstrap does not resample common date-level dependence in the Stage 4a innovations. Such dependence is excluded by Assumption 3(ii).

Let $w_i^\star$ and $w_t^\star$ denote the firm and date multipliers. Stage 1 is the weighted cross-sectional

regression

$$\widehat{a}_t^\star = \left(\sum_{i \in \mathcal{T}_t^{(0)}} w_i^\star Z_{i,t-1} Z_{i,t-1}' \right)^{-1} \sum_{i \in \mathcal{T}_t^{(0)}} w_i^\star Z_{i,t-1} R_{it}.$$

Define $\widehat{v}_t = \widehat{a}_t - \bar{\widehat{a}} - \widehat{\delta}(D_{m,t-1} - \bar{D}_m)$. Stage 2 is rerun using

$$\widehat{y}_t^\star = \widehat{a}_t^\star + w_t^\star \widehat{B}\widehat{B}'\widehat{v}_t.$$

The date multiplier enters only this calendar regression and affects only $\widehat{\delta}^\star$.

Stage 3 is rerun using

$$\widehat{a}_t^{0,\star} = \widehat{a}_t^\star - \widehat{\delta}^{\star,N} D_{m,t-1},$$

where $\widehat{\delta}^{\star,N}$ is computed from the firm-weighted slopes $\widehat{a}_t^\star$ without the date multiplier. This choice propagates the Stage 1 draw into PCA without adding the calendar residual. The residualization must use the firm-only coefficient. The proof of Theorem 2 shows that using $\widehat{\delta}^\star$ would eliminate the exact calendar annihilation and introduce a divergent term into the loading draw.

Stage 4a is the $w_i^\star$-weighted regression of $R_{it} - Z_{i,t-1}'\widehat{a}_t^\star$ on $(Z_{i,t-1}', (Z_{i,t-1} \otimes \widehat{F}_t^\star)')'$. Define

$$H^\star = (\widehat{F}' M_T \widehat{F}^\star)(\widehat{F}^{\star\prime} M_T \widehat{F}^\star)^{-1}.$$

The loading draws are centered at $\widehat{B}H^\star$, and the exposure draws are centered at $\widehat{\Gamma}H^\star$. Only the discrete column signs are aligned. Accordingly, $\widehat{\vartheta}^\star - \widehat{\vartheta}$ denotes the rotation-centered deviations for each parameter block.

Remark 3 (Path-mean augmentation in the bootstrap). The average effects in Remark 2 depend on two factor-path means that are not included in ϑ. Each bootstrap replication redraws these means using the same date multiplier $w_t^\star$ as the Stage 2 calendar draw. Define

$$W^\star = \frac{1}{T_0} \sum_{t:\, D_{m,t-1}=0} w_t^\star (\widehat{F}_t^\star - \bar{\widehat{F}}_0^\star), \qquad W_g^\star = \frac{1}{T_g} \sum_{t:\, D_{m,t-1}^{(g)}=1} w_t^\star (\widehat{F}_t^\star - \bar{\widehat{F}}_g^\star).$$

Here, $\widehat{F}^\star$ is the factor path in the bootstrap draw. The quantities $\bar{\widehat{F}}_0^\star$ and $\bar{\widehat{F}}_g^\star$ are its means over release-free dates and type-g announcement dates. The corresponding numbers of dates are T_0 and T_g. For the clean effects, the bootstrap replaces $\widehat{\mu}^\star$ with $\widehat{\mu}^\star + W^\star$. For the type-$g$ announcement-day effects, it uses $\widehat{\mu}^\star + W_g^\star$.[11]

Under Assumption 1(i), these wild-bootstrap averages reproduce the $\sqrt{T}$-scale distributions of the corresponding path means. Using the same date multiplier also preserves their covariance with the calendar score. Corollary 2 gives the result for the clean effect, and Proposition 3(iii) gives the

[11] The two displays presume that no date carries more than one macro release, so that the projection of the factor path on the calendars is saturated and its intercept and type-g fitted value are the two group averages. When calendars overlap, $W^\star$ and $W_g^\star$ are instead the intercept and the type-g fitted value of a wild-multiplier least-squares fit of $\widehat{F}_t^\star$ on $(1, \boldsymbol{D}_{t-1}^{m\prime})'$. In the sample of Section 4.1, 61 of 11,323 dates carry two releases.

result for the announcement-day effect.

Assumption 5 (Bootstrap weights and influence regularity)**.**

(i) Multipliers. *The firm multipliers* $\{w_i^\star\}_{i\le N}$ *and date multipliers* $\{w_t^\star\}_{t\le T}$ *are i.i.d. within each array, mutually independent, and independent of the data, with*

$$\mathbb{E}w_i^\star = 1, \quad \mathrm{Var}(w_i^\star) = 1, \qquad \mathbb{E}w_t^\star = 0, \quad \mathrm{Var}(w_t^\star) = 1,$$

and $\mathbb{E}|w_i^\star|^{4+\varepsilon} + \mathbb{E}|w_t^\star|^{4+\varepsilon} < \infty$ *for some* $\varepsilon > 0$*.*

(ii) Feasible influence arrays. *Let* $\phi_{t,NT}$ *and* $\psi_{i,NT}$ *denote the normalized date and firm contributions to the joint linear expansion*

$$\mathcal{R}_{NT}(\widehat{\vartheta} - \vartheta) = \sum_{t=1}^{T} \phi_{t,NT} + \sum_{i=1}^{N} \psi_{i,NT} + o_p(1).$$

Let $\widehat{\phi}_{t,NT}$ *and* $\widehat{\psi}_{i,NT}$ *denote their feasible residual and sample-Jacobian counterparts used by the bootstrap. These arrays satisfy*

$$\sum_{t=1}^{T} \left\|\widehat{\phi}_{t,NT} - \phi_{t,NT}\right\|^2 + \sum_{i=1}^{N} \left\|\widehat{\psi}_{i,NT} - \psi_{i,NT}\right\|^2 \xrightarrow{p} 0$$

and

$$\max_{t\le T} \left\|\widehat{\phi}_{t,NT}\right\| + \max_{i\le N} \left\|\widehat{\psi}_{i,NT}\right\| \xrightarrow{p} 0.$$

The same conditions hold for the fixed-dimensional stacked arrays. The infeasible outer products satisfy

$$\sum_{t=1}^{T} \phi_{t,NT}\phi_{t,NT}' + \sum_{i=1}^{N} \psi_{i,NT}\psi_{i,NT}' \xrightarrow{p} \mathcal{V},$$

where $\mathcal{V}$ *is the covariance matrix defined in* (26)*.*

(iii) Cross-array orthogonality. *With* $\phi_{t,NT}$ *and* $\psi_{i,NT}$ *as in part (ii),*

$$\mathrm{Cov}\left(\sum_{t\le T} \phi_{t,NT}, \sum_{i\le N} \psi_{i,NT}\right) \longrightarrow 0.$$

Assumption 5 gives the additional conditions needed for the split bootstrap. Part (i) specifies independent firm and date multipliers. The i.i.d. mean-zero date multiplier is appropriate because the date-indexed scores are martingale differences under Assumptions 1(i) and 3(ii). Serial dependence in these scores would require block or cluster multipliers. The firm multiplier relies on the maintained conditional independence of innovations across firms. Cross-firm clustering would require cluster-level multipliers and corresponding changes to the influence arrays.

Assumption 4(ii) provides the joint central limit theorem for the infeasible score array. Part (ii) of Assumption 5 instead requires the feasible influence arrays to approximate their infeasible counterparts. It also requires their largest contributions to vanish and their outer products to converge to $\mathcal{V}$. These conditions ensure that the feasible bootstrap reproduces the covariance of the limiting score array.

Independent date and firm multiplier arrays reproduce the separate covariances of the two influence sums and impose zero covariance between them. The orthogonality condition in Assumption 5(iii) is therefore needed for the split bootstrap to recover $\mathcal{V}$. The maintained conditional structure motivates this restriction. The date-indexed array is driven by factor and calendar innovations that are mean zero given $\mathcal{G}_{t-1}$, while the firm-indexed array is driven by conditionally mean-zero return innovations. We nevertheless impose orthogonality directly. The condition does not restrict the cross-stage parameter covariances contained within the firm-indexed influence sum.

Deriving the conditions in Assumption 5(ii) and (iii) from lower-level sampling and moment restrictions would require a conditional Lindeberg argument of the type established by Gonçalves and Perron (2014) under the multiplier moment conditions of Mammen (1993). The firm and date multipliers act on different stages of the sequential estimator, so their joint conditional limit must be established. We impose the required influence-array conditions directly. Theorem 2 is conditional on these requirements. Internet Appendix A provides finite-sample evidence for the resulting approximation.

The bootstrap also conditions on the chosen factor dimension. Each draw extracts the same K factors, so the reported standard errors do not account for uncertainty in the factor count.

Theorem 2 (Bootstrap validity)**.** *Suppose Assumptions 1, 2, 3, 4, and 5 hold and $T/N \to 0$. Then, conditionally on the sample,*

$$\sup_{f \in \mathrm{BL}_1(\mathbb{R}^p)} \left| \mathbb{E}^\star f\left\{ \mathcal{R}_{NT}(\widehat{\vartheta}^\star - \widehat{\vartheta}) \right\} - \mathbb{E} f\left\{ \mathcal{R}_{NT}(\widehat{\vartheta} - \vartheta) \right\} \right| \xrightarrow{p} 0,$$

where $p := \dim \vartheta = L(2K+3)$ and $\mathrm{BL}_1(\mathbb{R}^p)$ is the class of functions bounded in absolute value by one with Lipschitz constant at most one; equivalently, the bootstrap law converges weakly, conditionally on the sample and in probability, to $\mathcal{N}(0, \mathcal{V})$, the same law as $\mathcal{R}_{NT}(\widehat{\vartheta} - \vartheta)$. [12]

Corollary 1 (Bootstrap inference on smooth functionals)**.** *Let $g : \mathbb{R}^p \to \mathbb{R}$ and put*

$$\sigma_{NT}^2 := \nabla g(\vartheta)' \mathcal{R}_{NT}^{-1} \mathcal{V} \mathcal{R}_{NT}^{-1} \nabla g(\vartheta).$$

Suppose the assumptions of Theorem 2 hold and, in addition:

(a) g is continuously differentiable in a neighborhood of ϑ, and $\sigma_{NT} > 0$ for all large (N, T).

[12] The bounded-Lipschitz metric is the standard metrization of conditional weak convergence in probability and is what Corollary 1 requires. Because $\mathcal{V}$ is singular the limit places its mass on a proper subspace of $\mathbb{R}^p$ and has no Lebesgue density.

(b) Whenever the leading term of σ_{NT} is of order faster than $T^{-1/2}$, the δ block of $\nabla g(\vartheta)$ lies in $\mathrm{col}(B)$.

(c) g admits the rate-adapted expansion

$$\sup_{\|h\|\le M}\left|g(\vartheta+\mathcal{R}_{NT}^{-1}h)-g(\vartheta)-\nabla g(\vartheta)'\mathcal{R}_{NT}^{-1}h\right|=o(\sigma_{NT})\quad \textit{for every fixed } M<\infty.$$

Then

$$\sup_{x\in\mathbb{R}}\left|\mathbb{P}^\star\Big\{\sigma_{NT}^{-1}\{g(\widehat{\vartheta}^\star)-g(\widehat{\vartheta})\}\le x\Big\}-\mathbb{P}\Big\{\sigma_{NT}^{-1}\{g(\widehat{\vartheta})-g(\vartheta)\}\le x\Big\}\right|\xrightarrow{p}0,$$

so percentile and studentized bootstrap intervals for $g(\vartheta)$ have asymptotically correct coverage and the bootstrap variance of $g(\widehat{\vartheta})$ is consistent for σ_{NT}^2.

Conditions (b) and (c) address the two ways in which σ_{NT} can fail to describe the fluctuation of $g(\widehat{\vartheta})$. The matrix $\mathcal{V}$ is degenerate in the $\ker(B')$ directions of the δ block and in the K joint (α, B) directions of (27), but the two degeneracies are different. The second is induced by the exact sample restriction $\widehat{\alpha}'\widehat{B}=0$, whose first-order linearization vanishes asymptotically. The first is a consequence of the $\sqrt{T}$ scaling. By (25) the span residual $(I_L-BB')\hat{\delta}$ still fluctuates at the $(NT)^{-1/2}$ order, so a $\ker(B')$ component of the δ gradient can contribute at the same order as an $(NT)^{-1/2}$-rate leading term while leaving σ_{NT}^2 unchanged. Condition (b) removes that case. Condition (c) is needed because the convergence is at mixed rates. A functional with zero gradient but nonzero curvature in a $\sqrt{T}$ direction has a second-order term of order T^{-1}, which under $T/N\to 0$ dominates an $(NT)^{-1/2}$ first-order standard error, the ratio growing like $\sqrt{N/T}$.

Every functional reported through this corollary satisfies both conditions, for one of two trivial reasons. For $z'\alpha$, $z'\gamma$, and any smooth function of BB', the δ gradient is zero, so (b) holds and the curvature lies in the $\sqrt{NT}$ blocks, making the remainder in (c) of order $(NT)^{-1}$. For $z'\delta$ the δ gradient is z, but $\sigma_{NT}^2=z'V_\delta z/T$, so the leading term is of order $T^{-1/2}$ and (b) does not bind, while the functional is affine and the remainder in (c) vanishes identically. The span share $\delta'BB'\delta/\delta'\delta$ is not covered at $\delta=B\lambda$ as its first-order derivative vanishes, so $\sigma_{NT}=0$ and condition (a) fails. We report it as a descriptive decomposition and attach no interval to it. Functions of BB' are covered even though the parameter block is $\widehat{B}H^{-1}$, which is not orthonormal, because $\widehat{B}\widehat{B}'$ is the projector map applied to $\widehat{B}H^{-1}$, a smooth function of ϑ with population value BB'. The average firm-event effect replaces the baseline mean by $\widehat{\mu}$ and is not a functional of ϑ alone, so it is treated in Corollary 2.

Corollary 2 (The average event effect)**.** *Fix $z\in\mathbb{R}^L$ and a reported firm-event type j. Define the estimator, its estimand, and its bootstrap draw by*

$$\widehat{c}_z:=z'(\widehat{\gamma}_j+\widehat{\Gamma}_j\widehat{\mu}),\qquad c_z:=\tau^\gamma(z)=z'(\gamma_j+\Gamma_j\mu),\qquad \widehat{c}_z^\star:=z'(\widehat{\gamma}_j^\star+\widehat{\Gamma}_j^\star(\widehat{\mu}^\star+W^\star)),$$

where every draw recomputes $\widehat{\mu}^\star:=T^{-1}\sum_t\widehat{F}_t^\star$ and the path-mean redraw $W^\star$ of Remark 3. Suppose the assumptions of Theorem 2 hold and, in addition:

(a) The joint central limit theorem of Assumption 4(ii) is read as including the in-span $\sqrt{NT}$-scaled component of the calendar-noise score together with the $\sqrt{T}$-scaled path score $T^{-1/2}\sum_t(\bar{F}_t-\mu)$.

(b) Let $\sigma^2_{c,NT}$ be the variance at sample size (N,T) of the sum of four first-order terms, namely the three terms of the exact decomposition (C.42) and the path term $z'\Gamma_j(\mu_T^0-\mu)$. With $n_1^{(j)}$ the type-j clean-cell size of Assumption 3(iii), the scaled limit $\sigma_c^2 := \lim (n_1^{(j)} \wedge T)\,\sigma^2_{c,NT}$ exists and is positive.

Then

$$\sigma_{c,NT}^{-1}(\widehat{c}_z - c_z) \Rightarrow \mathcal{N}(0,1), \qquad \sup_{x\in\mathbb{R}} |\mathbb{P}^\star\{\sigma_{c,NT}^{-1}(\widehat{c}_z^\star - \widehat{c}_z) \le x\} - \mathbb{P}\{\sigma_{c,NT}^{-1}(\widehat{c}_z - c_z)\le x\}| \xrightarrow{p} 0,$$

so percentile and studentized intervals for c_z have asymptotically correct coverage.

It is Corollaries 1 and 2, and not Theorem 2, that license the reported standard errors, since the joint limit is degenerate while every quantity reported with an interval is a scalar functional with nondegenerate limiting variance. Fitted conditional returns at a specific date fall outside both, because they involve the date-level factor estimate, and carry no intervals. The normalization in Corollary 2 is T, since $n_1^{(j)} \asymp NT$, so whenever $z'\Gamma_j \neq 0$ the average event effect is learned at the rate of the factor, and the width of its interval is set by how precisely the baseline factor mean is estimated.

B.4 Additional Inference Results

The aspect-ratio condition $T/N \to 0$ of Theorem 1 binds only the uncorrected estimator; it is consumed by exactly one feature of the proof: the systematic $O(N^{-1})$ generated-regressor and eigenvector biases of (C.17), (C.20), and (C.28), which contribute $\sqrt{T/N}$ after $\sqrt{NT}$ or $\sqrt{n_1}$ scaling. Every stochastic remainder in the proof is controlled whenever $T = o(N^3)$. A cross-sectional jackknife therefore weakens the condition for the corrected blocks.

Proposition 2 (Fixed-aspect-ratio jackknife)**.** *For $s = 1,\dots,S$ with S fixed, let $(\mathcal{N}_{(s,1)},\mathcal{N}_{(s,2)})$ be independent random partitions of the firms into halves, drawn independently of the data, and let $\widehat{\vartheta}_{(s,h)}$ be the four-stage estimates computed on the corresponding half cross-section. Align each half's loading and exposure blocks to the full-sample coordinates by the regression alignment of (22), computed with the half-sample factor path in the population slot and the full-sample path in the estimated slot. Define*

$$\widehat{\vartheta}^{\mathrm{jk}} := 2\,\widehat{\vartheta} - \frac{1}{2S}\sum_{s=1}^{S}(\widehat{\vartheta}_{(s,1)} + \widehat{\vartheta}_{(s,2)}).$$

Suppose Assumptions 1–4 hold and, in addition:

(a) Split-stable biases. *Write n for the cross-sectional size of the estimating sample, so that $n = N$*

for the full panel and $n = N/2$ for each half. The full-sample bias terms of (C.17), (C.20) and (C.28), and the bias of each aligned half-sample estimator, admit expansions $b\,n^{-1}+O_p(n^{-3/2})$ with the same coefficient b. The half-sample bias combines those same three terms at the half cross-section with the $O_p(n^{-1})$ noise-Gram term of the feasible alignment.

(b) $T/N^2 \to 0$.

Then:

(i) $\mathcal{R}_{NT}(\widehat{\vartheta}^{\mathrm{jk}} - \vartheta) \;\Rightarrow\; \mathcal{N}(0, \mathcal{V})$, *with the covariance* $\mathcal{V}$ *of* (24).

(ii) Under Assumption 5, the full-sample bootstrap standard errors of Theorem 2 remain consistent for the jackknifed blocks, and intervals centered at $\widehat{\vartheta}^{\mathrm{jk}}$ *with those standard errors have asymptotically correct coverage.*

With a fixed number of event types the statement applies blockwise to the stacked vector, with the type-j clean-cell sizes of Assumption 3(iii).

No component of the limit in (i) requires $T/N \to 0$. The choice of S affects only higher-order noise, and nothing is claimed as $S \to \infty$. The bootstrap clause of Proposition 2 carries this validity to intervals centered at the jackknifed estimator under $T/N^2 \to 0$, given the split-stability condition stated there.

The coincidence pair is excluded from the stacked vector for a reason, not by neglect. A direct expansion of the Stage-4b normal equations shows that the level block $\widehat{c}_g$ carries the Stage-2 calendar variation as an $O_p(T^{-1/2})$ common translation inside $\mathrm{col}(\Gamma_a)$, entangled with its own cell-level noise in the same coordinates: scaled by the coincidence cell's own root sample size the translation diverges, while at the $\sqrt{T}$ scaling the block only replicates the δ block's law. More decisively, the split bootstrap cannot reproduce the translation at all, the date multiplier enters Stage 2 alone, so Theorem 2 would fail for any vector containing $\widehat{c}_g$. The stacked vector is thus the largest vector of raw stage-coefficient blocks for which the joint limit and the bootstrap result both hold; the coincidence cell lies outside it. The reported coincidence quantities are instead the translation-cancelled $\hat{\gamma}_{a,g}$ and the announcement-day average effect, covered by Proposition 3; $\widehat{\Gamma}_a$ is displayed descriptively. The factor announcement premium needs no coordinate of its own: $\hat{B}\hat{\boldsymbol{\lambda}} = \widehat{P}_B\widehat{\boldsymbol{\delta}}$ is an exact function of the four-stage estimates, covered by Corollary 1.

Proposition 3 (Coincidence-cell estimation)**.** *Let $\mathcal{T}^{(1,m)}$ be the coincidence cell of (16), $n_{1m} := |\mathcal{T}^{(1,m)}|$, $n_{1m,t}$ the number of its observations on date t, and $n_{1m}^{(g)} := \#\{(i,t) \in \mathcal{T}^{(1,m)} : D_{m,t-1}^{(g)} = 1\}$. Let (c_g, Γ_a) be the Stage-4b pooled projection coefficients of Section 3.2 and $\gamma_{a,g} := c_g - \Gamma_a\lambda_g$. With H of (22), write the aligned blocks*

$$\check{\lambda}_g := H\hat{\lambda}_g, \qquad \check{\Gamma}_a := \hat{\Gamma}_a H^{-1},$$

centered at λ_g and Γ_a.[13] *Suppose Assumptions 1–4 hold with $T/N \to 0$ and, in addition:*

(a) $n_{1m}/(NT) \to \pi_{1m} \in (0,1)$, $n_{1m}^{(g)}/n_{1m} \to \pi^{(g)} \in (0,1)$ *for every reported macro type g, and* $\max_{t \le T} n_{1m,t}/n_{1m} \to 0$*;*

(b) the probability limit of the Stage-4b design second moment exists and is nonsingular;

(c) the joint central limit theorem of Assumption 4(ii) holds for the influence array of Stages 1–4b, and its feasible split-weighted version is square-consistent as in Assumption 5(ii).

Then, reading all factor-space objects in the coordinates of Theorem 1:

(i)

$$\hat{c}_g - c_g = \Gamma_a B'(\hat{\delta}_g - \delta_g) + r_{c,g}, \qquad \check{\lambda}_g - \lambda_g = B'(\hat{\delta}_g - \delta_g) + r_{\lambda,g},$$

where the shared term $\Gamma_a B'(\hat{\delta}_g - \delta_g)$ is $O_p(T^{-1/2})$ and $\sqrt{n_{1m}}(r_{c,g},\, r_{\lambda,g},\, \mathrm{vec}(\check{\Gamma}_a - \Gamma_a))$ is jointly asymptotically Gaussian, jointly over the reported macro types.

(ii) With $\hat{\gamma}_{a,g} = \hat{c}_g - \hat{\Gamma}_a \hat{\lambda}_g$,

$$\hat{\gamma}_{a,g} - \gamma_{a,g} = r_{c,g} - \Gamma_a r_{\lambda,g} - (\check{\Gamma}_a - \Gamma_a)\lambda_g + o_p(n_{1m}^{-1/2}),$$

so the $O_p(T^{-1/2})$ term of (i) cancels, and $\sqrt{n_{1m}}(\hat{\gamma}_{a,g} - \gamma_{a,g}, \mathrm{vec}(\check{\Gamma}_a - \Gamma_a))$ is jointly asymptotically normal, jointly over the reported macro types.

(iii) Suppose every bootstrap draw reruns Stages 1–4b and forms, within that draw,

$$\hat{\lambda}_g^{\star,N} := (\hat{B}^{\star\prime}\hat{B}^{\star})^{-1}\hat{B}^{\star\prime}\hat{\delta}_g^{\star,N}, \qquad \hat{\gamma}_{a,g}^{\star} = \hat{c}_g^{\star} - \hat{\Gamma}_a^{\star}\hat{\lambda}_g^{\star,N}, \qquad z'(\hat{c}_g^{\star} + \hat{\Gamma}_a^{\star}(\hat{\mu}^{\star} + W_g^{\star})),$$

where $\hat{\delta}_g^{\star,N}$ is the firm-only Stage-2 coefficient used to residualize Stage 3 and $W_g^{\star}$ is the type-g path-mean redraw of Remark 3. Then, conditionally on the sample and in probability, the law of the centered vector converges to the joint limit of (ii), and the effect draw covers $\tau^{\gamma\delta}(z) = z'(c_g + \Gamma_a \mu)$ under the augmented central limit theorem of Corollary 2.

The bootstrap covariance in (iii) retains the cross covariances that produce the cancellation in (ii), which combining a Stage-4b variance for $\hat{c}_g$ with a separate variance for $\hat{\lambda}_g$ would lose. The reduced-form intercept $\hat{c}_g$ alone retains the mixed-rate calendar component of (i) and stays outside the stacked vector.

C Proofs

Throughout the proofs $P_B := B(B'B)^{-1}B'$ is the projector onto $\mathrm{col}(B)$ and a_t is the date-specific characteristic slope of (17). For a full-column-rank Γ_a write $P_{\Gamma_a} := \Gamma_a(\Gamma_a'\Gamma_a)^{-1}\Gamma_a'$, and similarly

[13] Their product is alignment-free, $\check{\Gamma}_a\check{\lambda}_g = \hat{\Gamma}_a\hat{\lambda}_g$, so $\hat{\gamma}_{a,g} := \hat{c}_g - \hat{\Gamma}_a\hat{\lambda}_g$ does not depend on the convention.

$P_{\widehat{\Gamma}_a}$; note $P_{\Gamma_a H} = P_{\Gamma_a}$ for every nonsingular H. For a firm event let

$$q_{it} := R_{it} - Z'_{i,t-1} a_t, \qquad x_{it} := \begin{pmatrix} Z_{i,t-1} \\ Z_{i,t-1} \otimes \bar{F}_t \end{pmatrix},$$

and on a firm–macro coincidence date let c and Γ_a be the reduced-form coefficients in

$$\mathbb{E}[q_{it} \mid \mathcal{I}_{t-1}, \bar{F}_t, D_{i,t-1} = D_{m,t-1} = 1] = Z'_{i,t-1} c + Z'_{i,t-1} \Gamma_a \bar{F}_t,$$

so that c, not γ_a, is the object the coincidence cell delivers directly. Let $M := M_T$, $d_t := D_{m,t-1}$, $\widetilde{d}_t := d_t - \bar{D}_m$, and $\widehat{v}_m := T^{-1} \sum_t \widetilde{d}_t^2$. Slopes are columns: $A := (a_1, \ldots, a_T) \in \mathbb{R}^{L \times T}$, while $\bar{F} \in \mathbb{R}^{T \times K}$ has row t equal to $\bar{F}'_t$. Let

$$\widehat{Q}_{0t} := \frac{1}{N_t^{(0)}} \sum_{i: D_{i,t-1}=0} Z_{i,t-1} Z'_{i,t-1}, \qquad \widetilde{\epsilon}_t := \widehat{Q}_{0t}^{-1} \frac{1}{N_t^{(0)}} \sum_{i: D_{i,t-1}=0} Z_{i,t-1} U_{it}(0). \tag{C.1}$$

The proof uses N as the cross-sectional rate index; the nonvanishing control fraction of Assumption 3(i) makes this equivalent to using $\min_t N_t^{(0)}$.

C.1 Proof of Proposition 1

Proof of Proposition 1. *(i)* Fix a date t. The slope $a_t = \alpha + B\bar{F}_t + \delta d_t$ is $(\mathcal{G}_{t-1}, \bar{F}_t)$-measurable, so for a control firm Assumption 1(i) gives $\mathbb{E}[R_{it} \mid \mathcal{G}_{t-1}, \bar{F}_t] = Z'_{i,t-1} a_t$. Averaging these conditional moments over the control cross-section,

$$m_{0t} := \frac{1}{N_t^{(0)}} \sum_{i: D_{i,t-1}=0} Z_{i,t-1}\, \mathbb{E}[R_{it} \mid \mathcal{G}_{t-1}, \bar{F}_t] = \widehat{Q}_{0t}\, a_t, \qquad a_t = \widehat{Q}_{0t}^{-1} m_{0t}, \tag{C.2}$$

the Gram $\widehat{Q}_{0t}$ of (C.1) being nonsingular almost surely by Assumption 1(ii). Conditioning (C.2) down to $\mathcal{G}_{t-1}$ and using $\mathbb{E}[\bar{F}_t \mid \mathcal{G}_{t-1}] = \mu$,

$$\mathbb{E}[a_t \mid \mathcal{G}_{t-1}] = \alpha + B\mu + \delta\, d_t. \tag{C.3}$$

Thus δ has the property that $\mathbb{E}[a_t \mid \mathcal{G}_{t-1}] - \delta d_t$ is a single nonrandom vector common to all t, and it is the only vector with that property: if δ' also satisfies the same property, subtracting the two statements gives $(\delta - \delta')d_t = k$ for a fixed k, for every t and almost surely, and by Assumption 1(iii) the binary d_t equals 0 on a set of positive probability, forcing $k = 0$, and equals 1 on a set of positive probability, forcing $\delta = \delta'$. Hence δ is identified, and with it

$$a_t^0 := a_t - \delta d_t = \alpha + B\bar{F}_t, \qquad \bar{a} := \mathbb{E}[a_t^0 \mid \mathcal{G}_{t-1}] = \alpha + B\mu, \qquad a_t^0 - \bar{a} = B(\bar{F}_t - \mu).$$

The centered second moments of the identified path are finite by Assumption 1(i), so

$$S_0 := \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}\big[(a_t^0 - \bar{a})(a_t^0 - \bar{a})'\big] = B\Sigma_F B' \tag{C.4}$$

is identified, with Σ_F the limit of Assumption 2(i). Since $S_0 x = B\Sigma_F B' x$ for every x, $\mathrm{col}(S_0) \subseteq \mathrm{col}(B)$; and $\Sigma_F \succ 0$ with $\operatorname{rank} B = K$ give $\operatorname{rank} S_0 = K$, so the two spaces coincide. Hence P_B is identified. Fix an $L \times K$ matrix B_0 with orthonormal columns spanning $\mathrm{col}(S_0)$ and define

$$g_t := B_0'(a_t^0 - \bar{a}) = R(\bar{F}_t - \mu), \qquad R := B_0' B,$$

identified date by date, with R nonsingular because $B = B_0 R$ has rank K.

On the clean-event cell, (11) and Assumption 1(i) give, $D_{i,t-1}$ being $\mathcal{G}_{t-1}$-measurable,

$$\mathbb{E}[q_{it} \mid Z_{i,t-1}, \bar{F}_t, D_{i,t-1} = 1, D_{m,t-1} = 0] = x_{it}'\theta_\gamma^0, \qquad \theta_\gamma^0 = \begin{pmatrix} \gamma \\ \mathrm{vec}(\Gamma') \end{pmatrix}. \tag{C.5}$$

The design x_{it} involves the unidentified $\bar{F}_t$, but

$$Z_{i,t-1}'(\gamma + \Gamma\bar{F}_t) = Z_{i,t-1}'(\gamma + \Gamma\mu) + (Z_{i,t-1} \otimes g_t)' \,\mathrm{vec}\,((\Gamma R^{-1})')$$

holds identically in $(Z_{i,t-1}, \bar{F}_t)$, so (C.5) may be read on the identified design $w_{it} := (Z_{i,t-1}', (Z_{i,t-1} \otimes g_t)')'$. The two designs are linked by the nonsingular constant map $w_{it} = \mathcal{T} x_{it}$, $\mathcal{T} := \begin{pmatrix} I_L & 0 \\ -(I_L \otimes R\mu) & I_L \otimes R \end{pmatrix}$, so the pooled clean-cell Gram of w_{it} is nonsingular by Assumption 1(iii). Because the displayed rewriting is an identity and not a projection, the unique solution of $\mathbb{E}[w_{it}(q_{it} - w_{it}'\theta)] = 0$ has intercept block $\gamma + \Gamma\mu$ and slope block $\Pi := \Gamma R^{-1}$, both therefore identified. With them $\Gamma(\bar{F}_t - \mu) = \Pi g_t$, and every fitted conditional return outside the coincidence cell: $Z_{i,t-1}'a_t$ for a control firm and $Z_{i,t-1}'a_t + Z_{i,t-1}'(\gamma + \Gamma\mu) + Z_{i,t-1}'\Pi g_t$ for a clean event firm. Replacing B_0 by $B_0 O$ for orthogonal O replaces g_t by $O'g_t$ and Π by $\Pi O'^{-1}$, leaving Πg_t unchanged.

Finally, for any $\zeta \in \mathbb{R}^K$ the transformation

$$\bar{F}_t^\dagger = \bar{F}_t - \zeta, \qquad \alpha^\dagger = \alpha + B\zeta, \qquad \gamma^\dagger = \gamma + \Gamma\zeta$$

leaves $\alpha^\dagger + B\bar{F}_t^\dagger = \alpha + B\bar{F}_t$ and $\gamma^\dagger + \Gamma\bar{F}_t^\dagger = \gamma + \Gamma\bar{F}_t$, hence every conditional mean in the model, unchanged, and satisfies $\mathbb{E}[\bar{F}_t^\dagger \mid \mathcal{G}_{t-1}] = \mu - \zeta$. Taking $\zeta \neq 0$ moves α and μ, and moves γ whenever $\Gamma \neq 0$, while fixing $\alpha + B\mu$ and $\gamma + \Gamma\mu$. So the three are not separately identified.

(ii) Assumption 2(ii) states $\alpha' B = 0$, that is $P_B\alpha = 0$. Applying P_B and $I_L - P_B$ to $\bar{a} = \alpha + B\mu$ separates the two terms,

$$\alpha = (I_L - P_B)\bar{a}, \qquad B\mu = P_B\bar{a},$$

both identified by (i). Then $R\mu = B_0'B\mu = B_0'P_B\bar{a}$ is identified, and with it $\Gamma\mu = \Pi(R\mu)$ and

$\gamma = (\gamma + \Gamma\mu) - \Gamma\mu$. Adding the centered components of (i) gives $B\bar{F}_t = B(\bar{F}_t - \mu) + B\mu$ and $\Gamma\bar{F}_t = \Pi g_t + \Gamma\mu$. Each of these objects is a projection or a product of the form $B\cdot$ or $\Gamma\cdot$, so the reparametrization $B \mapsto BQ^{-1}$, $\bar{F}_t \mapsto Q\bar{F}_t$, $\Gamma \mapsto \Gamma Q^{-1}$ leaves each unchanged, so none depends on the factor coordinate system.

(iii) Assumption 2(iii) adds $B'B = I_K$ and a diagonal Σ_F with strictly decreasing diagonal entries. Under these, the right-hand side of (C.4) is the spectral decomposition of the identified matrix S_0. Its K nonzero eigenvalues are the diagonal entries of Σ_F and the associated eigenvectors are the columns of B. Distinctness of those entries makes each eigenvector unique up to its sign, so B is identified up to a sign in each column. Consequently $R = B_0'B$ is identified, and $\Gamma = \Pi R$. Premultiplying $a_t^0 = \alpha + B\bar{F}_t$ and $\bar{a} = \alpha + B\mu$ by B' and using $\alpha'B = 0$ and $B'B = I_K$,

$$\bar{F}_t = B'a_t^0, \qquad \mu = B'\bar{a},$$

so the level of the factor is recovered, not only $\bar{F}_t - \mu$. The residual indeterminacy is exactly the column signs: replacing B by BQ, $\bar{F}_t$ by $Q\bar{F}_t$ and Γ by ΓQ with $Q = \mathrm{diag}(\pm 1)$ reproduces S_0, $B\bar{F}_t$, $\Gamma\bar{F}_t$ and $\Gamma\mu$.

(iv) Under the hypothesis of part (iv), (11) and Assumption 1(i) give on the coincidence cell

$$\mathbb{E}[q_{it} \mid Z_{i,t-1}, \bar{F}_t, D_{i,t-1} = D_{m,t-1} = 1] = Z_{i,t-1}'c + Z_{i,t-1}'\Gamma_a\bar{F}_t, \tag{C.6}$$

and the reasoning in (C.5) in part (i) applies, with the coincidence-cell second moment in place of the clean-event one: it identifies $c + \Gamma_a\mu$ and $\Pi_a := \Gamma_a R^{-1}$, hence $\Gamma_a(\bar{F}_t - \mu) = \Pi_a g_t$. Part (ii) then gives $\Gamma_a\mu = \Pi_a(R\mu)$ and $c = (c + \Gamma_a\mu) - \Gamma_a\mu$, and part (iii) gives $\Gamma_a = \Pi_a R$.

(v) Under the factor-response model of Section 2.2, (10) gives $\delta = B\lambda$ and $F_t = \bar{F}_t + \lambda D_{m,t-1}$. Hence $BF_t = B\bar{F}_t + B\lambda D_{m,t-1} = B\bar{F}_t + \delta D_{m,t-1}$, which part (ii) identifies. Next, $R\lambda = B_0'B\lambda = B_0'\delta$ is identified by part (i), so $\Gamma_a\lambda = \Pi_a(R\lambda)$ is identified with Π_a from part (iv); since part (ii) also identifies c, the structural coincidence intercept $\gamma_a = c - \Gamma_a\lambda$ is identified. Adding $B'B = I_K$ from part (iii) gives $\lambda = B'\delta$ and therefore F_t itself. The product $\Gamma_a\lambda$ is invariant to the residual column signs, under which $\Gamma_a \mapsto \Gamma_a Q$ and $\lambda \mapsto Q\lambda$ with $Q = \mathrm{diag}(\pm 1)$, so γ_a does not depend on the sign convention.

On an announcement date, changing $F_t(0)$ and η_t of (5) by equal and opposite amounts leaves F_t, and hence every conditional mean in the model, unchanged. □

C.2 Proof of Theorem 1

Proof. The Stage-1 normal equations and (C.1) give the exact identity

$$\widehat{a}_t - a_t = \widetilde{\epsilon}_t. \tag{C.7}$$

By the uniform convergence and nonsingularity conditions in Assumption 3(i), $\min_t \lambda_{\min}(\widehat{Q}_{0t})$ is bounded away from zero with probability approaching one. On this event, the uniform Gram bounds, conditional cross-sectional independence, and conditional moment condition imply

$$\mathbb{E}[\|\widetilde{\epsilon}_t\|^2 \mid \mathcal{G}_{t-1}, \bar{F}_t] \leq \frac{C}{N_t^{(0)}}, \quad \frac{1}{T}\sum_t \|\widetilde{\epsilon}_t\|^2 = O_p(N^{-1}), \quad \left\|\frac{1}{T}\sum_t \widetilde{\epsilon}_t\right\| = O_p((NT)^{-1/2}). \tag{C.8}$$

The last statement uses the date martingale-difference property. Substitution of (C.7) into the Stage-2 normal equations gives

$$\widehat{\delta} - \delta = \widehat{v}_m^{-1}\frac{1}{T}\sum_t \widetilde{d}_t\{B(\bar{F}_t - \mu) + \widetilde{\epsilon}_t\}. \tag{C.9}$$

It follows that $\widehat{\delta} - \delta = O_p(T^{-1/2})$ and

$$\sqrt{T}(\widehat{\delta} - \delta) = v_m^{-1}T^{-1/2}\sum_t \widetilde{d}_t B(\bar{F}_t - \mu) + o_p(1), \tag{C.10}$$

because the $\widetilde{\epsilon}_t$ part is $O_p(N^{-1/2})$ at this scale.

The statistic in (C.10) is the combination

$$T^{-1/2}\sum_t d_t B(\bar{F}_t - \mu) - \bar{D}_m\, T^{-1/2}\sum_t B(\bar{F}_t - \mu)$$

of the two $\mathcal{G}_{t-1}$-adapted martingale-difference arrays of Assumption 3(ii). The joint central limit theorem applies to the adapted pair, and $\bar{D}_m \to_p \pi_m$ with $\widehat{v}_m \to_p v_m$ delivers the limit by Slutsky's theorem. The α block below uses the smaller orthogonal component,

$$\sqrt{NT}(I_L - BB')(\widehat{\delta} - \delta) = v_m^{-1}\sqrt{\frac{N}{T}}\sum_t \widetilde{d}_t(I_L - BB')\widetilde{\epsilon}_t + o_p(1). \tag{C.11}$$

Let $\widetilde{E} = (\widetilde{\epsilon}_1, \ldots, \widetilde{\epsilon}_T)$, $d = (d_1, \ldots, d_T)'$, and $\Delta^{(2)} = (\widehat{\delta} - \delta)d'$. The Stage-3 input obeys

$$\widehat{A}^0 = \alpha \mathbf{1}_T' + B\bar{F}' + \widetilde{E} - \Delta^{(2)}. \tag{C.12}$$

If V contains the K retained eigenvalues, the eigen-equation and $\widehat{F} = \widehat{A}^{0\prime}\widehat{B}$ imply

$$\widehat{B} - BH = \frac{1}{T}(\widetilde{E} - \Delta^{(2)})M\widehat{F}\,V^{-1}, \qquad H = (\bar{F}'M\widehat{F})(\widehat{F}'M\widehat{F})^{-1}. \tag{C.13}$$

The Stage-2 OLS normal equations are $\widehat{A}^0 M d = 0$, so

$$\Delta^{(2)}M\widehat{F} = (\widehat{\delta} - \delta)d'M\widehat{A}^{0\prime}\widehat{B} = 0,$$

and therefore

$$\widehat{B} - BH = \frac{1}{T}\widetilde{E}M\widehat{F}\,V^{-1} \quad \text{exactly.} \tag{C.14}$$

Assumption 2(i) defines Σ_F as the averaged population second moment. Assumption 3(ii) gives $T^{-1}\bar{F}'M\bar{F} = \Sigma_F + O_p(T^{-1/2})$, while Assumption 2(iii) requires Σ_F to be diagonal with strictly decreasing positive entries. Let $S_0 = B\Sigma_F B'$. Stage 3 imposes the corresponding sample normalization by requiring $\widehat{F}'M\widehat{F}/T$ to be diagonal. The population and sample diagonalizations differ by $O_p(T^{-1/2})$.

Write $P := B'\widehat{B}$ and $G := T^{-1}\widetilde{E}M\bar{F}$, so $\|G\| = O_p((NT)^{-1/2})$ by (C.8) and the date martingale-difference property. Substituting (C.12) into $\widehat{A}^0M\widehat{A}^{0\prime}/T$ and eliminating $\widehat{\delta} - \delta$ with (C.9) gives the exact decomposition

$$\frac{1}{T}\widehat{A}^0M\widehat{A}^{0\prime} = B\widetilde{\Sigma}_F B' + (BG' + GB') + \frac{1}{T}\widetilde{E}M\widetilde{E}' + R_S, \qquad \widetilde{\Sigma}_F := \frac{1}{T}\bar{F}'M_{[\mathbf{1}_T,d]}\bar{F}, \tag{C.15}$$

where $M_{[\mathbf{1}_T,d]}$ is the residual projector on the constant and the calendar, and

$$\|R_S\| = O_p(T^{-1/2}(NT)^{-1/2}) + O_p((NT)^{-1}).$$

The Stage-2 error enters (C.15) through $\widetilde{\Sigma}_F$ and through the cross terms between the factor path and $\widetilde{E}$ contained in R_S. The difference between the sample and population factor moments is

$$B\Delta_T B', \qquad \Delta_T := \widetilde{\Sigma}_F - \Sigma_F = O_p(T^{-1/2}).$$

This perturbation lies in $\mathrm{col}(B)$. Its ordered eigenvectors are therefore $B\widetilde{Q}$, where $\widetilde{Q}$ diagonalizes $\widetilde{\Sigma}_F$. The distinct diagonal entries of Σ_F imply $\widetilde{Q} = I_K + O_p(T^{-1/2})$. The perturbation thus changes only the alignment and is absorbed by H. It does not enter $\widehat{B} - BH$. Its cross terms with $\widetilde{E}$ are $O_p(T^{-1/2}(NT)^{-1/2}) = o_p((NT)^{-1/2})$.

Substituting $M\widehat{F} = M\bar{F}P + M\widetilde{E}'\widehat{B} - Md(\widehat{\delta} - \delta)'\widehat{B}$ into (C.14) gives

$$\widehat{B} - BH = GPV^{-1} + \Big(\frac{1}{T}\widetilde{E}M\widetilde{E}'\Big)\widehat{B}V^{-1} - \frac{1}{T}(\widetilde{E}Md)(\widehat{\delta} - \delta)'\widehat{B}V^{-1}. \tag{C.16}$$

Since $V \to_p \Sigma_F$ and $P \to_p I_K$ under the column-sign convention, with $\|PV^{-1} - \Sigma_F^{-1}\| = O_p(T^{-1/2})$, the composite influence map is the bounded linear map

$$\mathcal{L}_B(\Xi) := \Xi\,\Sigma_F^{-1}, \qquad \Xi \in \mathbb{R}^{L\times K},$$

acting on the $L \times K$ noise and factor cross moment. In this notation (C.16) reads

$$\widehat{B} - BH = \mathcal{L}_B\left(\frac{1}{T}\sum_t \widetilde{\epsilon}_t(\bar{F}_t - \mu)'\right) + b_{B,N} + R_{B,NT}, \tag{C.17}$$

where the quadratic-noise eigenvector bias is the conditional mean of the second term of (C.16),

$$b_{B,N} = \left(\frac{1}{T} \sum_{t=1}^{T} \frac{1}{N_t^{(0)}} \widehat{Q}_{0t}^{-1} \Omega_t \widehat{Q}_{0t}^{-1} \right) B \, \Sigma_F^{-1}, \quad \Omega_t := \frac{1}{N_t^{(0)}} \sum_{i: D_{i,t-1}=0} \mathbb{E}[Z_{i,t-1} Z_{i,t-1}' U_{it}(0)^2 \mid \mathcal{G}_{t-1}, \bar{F}_t],$$

so that $\|b_{B,N}\| = O_p(N^{-1})$, and

$$\|R_{B,NT}\| = o_p((NT)^{-1/2}) + o_p(N^{-1}). \tag{C.18}$$

Indeed, replacing $\bar{F}_t - T^{-1} \sum_s \bar{F}_s$ by $\bar{F}_t - \mu$ costs $O_p(T^{-1/2}(NT)^{-1/2})$, the centered part of the second term of (C.16) has conditional variance $O((N^2T)^{-1})$, the third term is $O_p(T^{-1/2}(NT)^{-1/2})$ because $T^{-1}\widetilde{E}Md = O_p((NT)^{-1/2})$ and $\widehat{\delta} - \delta = O_p(T^{-1/2})$, and the PV^{-1} replacement error is $O_p(T^{-1/2}(NT)^{-1/2})$. Thus

$$\|\widehat{B} - BH\| = O_p((NT)^{-1/2} + N^{-1}). \tag{C.19}$$

Under $T/N \to 0$ the bias in (C.17) is negligible after $\sqrt{NT}$ scaling, and the joint score central limit theorem in Assumption 4(ii) applied to its first term gives the loading block of the Gaussian limit.

The intercept is a smooth function of the same quantities, $\widehat{\alpha} = (I_L - \widehat{B}\widehat{B}')\bar{\widehat{a}}^0$. Expanding this map around $(B, \alpha + B\mu, \delta)$, using (C.8), (C.11) and (C.17), gives

$$\widehat{\alpha} - \alpha = (I_L - BB')\bar{\widetilde{\epsilon}} - \mathcal{D}P_B[\widehat{B} - BH](\alpha + B\mu) - \bar{D}_m(I_L - BB')(\widehat{\delta} - \delta) + b_{\alpha,N} + R_{\alpha,NT}, \tag{C.20}$$

where

$$\mathcal{D}P_B[G] = (I_L - BB')GB' + BG'(I_L - BB'), \qquad \|b_{\alpha,N}\| = O_p(N^{-1}),$$

and $R_{\alpha,NT} = o_p((NT)^{-1/2}) + o_p(N^{-1})$. Two $O_p(T^{-1/2})$ terms are larger than the $\sqrt{NT}$ target and both are removed by the projector. The span component of $\widehat{\delta} - \delta$, which is the $\sqrt{T}$ variation of (C.10) and lies in $\mathrm{col}(B)$, is annihilated by $I_L - BB'$, leaving the $\sqrt{NT}$ component of (C.11). Likewise $\bar{\widehat{a}}^0$ contains $B(T^{-1}\sum_s \bar{F}_s - \mu) = O_p(T^{-1/2})$ from expanding around $\alpha + B\mu$, and $(I_L - BB')B = 0$ annihilates it exactly. Its only surviving trace is the cross term with $\widehat{B} - BH$, of order $T^{-1/2}\{(NT)^{-1/2} + N^{-1}\} = o_p((NT)^{-1/2})$, which is part of $R_{\alpha,NT}$. Equations (C.11), (C.17) and (C.20) therefore give the $\sqrt{NT}$ intercept limit under $T/N \to 0$.

Put $W_B := \widehat{B} - BH$, $P := B'\widehat{B}$ and $\Psi := I_K - HP'$. On a clean, nonmacro Stage-4a date,

$$\widehat{F}_t - H^{-1}\bar{F}_t = \widehat{B}'\widetilde{\epsilon}_t + W_B'\alpha - H^{-1}\Psi\bar{F}_t, \tag{C.21}$$

which follows by substituting (C.12) into $\widehat{F}_t = \widehat{B}'\widehat{a}_t^0$. It separates the date-specific projection noise from two date-common loading and rotation errors.

Let

$$X_{it}^H := \begin{pmatrix} Z_{i,t-1} \\ Z_{i,t-1} \otimes H^{-1}\bar{F}_t \end{pmatrix}, \qquad \theta_\gamma^H := \begin{pmatrix} \gamma \\ \mathrm{vec}((\Gamma H)') \end{pmatrix},$$

and write $\widehat{X}_{it}$ for the actual Stage-4a design. The Stage-4a normal equations give the exact identity

$$\widehat{\theta}_\gamma - \theta_\gamma^H = \widehat{\Sigma}_4^{-1} \frac{1}{n_1} \sum_{(i,t)\in\mathcal{T}^{(1,0)}} \widehat{X}_{it}\{U_{it}(1) + r_{it}\}, \tag{C.22}$$

where

$$\widehat{\Sigma}_4 = \frac{1}{n_1}\sum \widehat{X}_{it}\widehat{X}_{it}', \qquad r_{it} = -Z_{i,t-1}'\widetilde{\epsilon}_t + Z_{i,t-1}'\Gamma(\bar{F}_t - H\widehat{F}_t). \tag{C.23}$$

By (C.8), (C.19) and (C.21),

$$\frac{1}{n_1}\sum \|\widehat{X}_{it} - X_{it}^H\|^2 = o_p(1), \qquad \widehat{\Sigma}_4 \to_p \Sigma_4. \tag{C.24}$$

The same bounds and conditional orthogonality imply

$$\frac{1}{\sqrt{n_1}}\sum (\widehat{X}_{it} - X_{it}^H)U_{it}(1) = o_p(1). \tag{C.25}$$

Products of the design error and r_{it} have conditional mean $b_{NT} = O_p(N^{-1})$, because both contain the same control-sample noise. To justify the centered replacement we need the two date-common errors of (C.21) as explicit influence sums. For W_B this is (C.17). For Ψ, note that $P = H + B'W_B$, so $\Psi = I_K - HH' - HW_B'B$, while $\widehat{B}'\widehat{B} = I_K$ gives $H'H - I_K = -(H'B'W_B + W_B'BH) - W_B'W_B = O_p(\|W_B\|)$. Rearranging these two identities, $\Psi = H(H'B'W_B + W_B'W_B)H^{-1}$ exactly. Since $H = I_K + O_p(T^{-1/2})$ under the column-sign convention, which is the in-span alignment noted in Remark 1, this gives $\Psi = B'W_B + O_p(T^{-1/2}\|W_B\|) + O_p(\|W_B\|^2)$, and both remainder terms are $o_p((NT)^{-1/2})$ under $T/N \to 0$, absorbed into R_Ψ below. Hence (C.17) may be written

$$W_B = \frac{1}{T}\sum_{s=1}^T \mathcal{L}_s + R_B, \qquad \Psi = \frac{1}{T}\sum_{s=1}^T \mathcal{P}_s + R_\Psi, \qquad \mathcal{L}_s := \widetilde{\epsilon}_s(\bar{F}_s - \mu)'\Sigma_F^{-1}, \quad \mathcal{P}_s := B'\mathcal{L}_s, \tag{C.26}$$

where $\mathcal{L}_s$ and $\mathcal{P}_s$ are date-s influence terms of conditional size $O_p(N^{-1/2})$, forming martingale difference arrays with respect to $\mathcal{G}_{s-1}$ because $\mathbb{E}[\widetilde{\epsilon}_s \mid \mathcal{G}_{s-1}, \bar{F}_s] = 0$, and $R_B, R_\Psi = O_p(N^{-1}) + o_p((NT)^{-1/2})$ by (C.18). Substitution of (C.21) and (C.26) into the product separates same-date diagonal terms, whose average conditional expectation is b_{NT}, from an off-diagonal degenerate martingale-array double sum. Bounded fourth moments and $\max_t n_{1t}/n_1 \to 0$ make the conditional second moment of the normalized off-diagonal sum $O(N^{-1})$, and the remainder terms are $o_p(1)$ by

(C.18). Thus the centered products are $o_p(1)$ after $n_1^{-1/2}$ aggregation, and

$$\begin{aligned}
\frac{1}{\sqrt{n_1}}\sum \widehat{X}_{it}\{U_{it}(1)+r_{it}\} &= S_{1,NT}+S_{2,NT}+S_{3,NT}+\sqrt{n_1}\,b_{NT}+o_p(1),\\
S_{1,NT} &:= \frac{1}{\sqrt{n_1}}\sum X_{it}^H U_{it}(1),\\
S_{2,NT} &:= -\frac{1}{\sqrt{n_1}}\sum_t W_t\widetilde{\epsilon}_t, \qquad W_t := \sum_{i:(i,t)\in\mathcal{T}^{(1,0)}} X_{it}^H Z'_{i,t-1},\\
S_{3,NT} &:= -\frac{1}{\sqrt{n_1}}\sum_t A_t(\widehat{F}_t - H^{-1}\bar{F}_t),\\
A_t &:= \sum_{i:(i,t)\in\mathcal{T}^{(1,0)}} X_{it}^H (Z'_{i,t-1}\Gamma H).
\end{aligned} \tag{C.27}$$

The second and third scores are generally correlated, because both are functions of the same control-firm innovations, and their cross covariance is part of $\mathcal{V}$.

Since $n_1 \asymp NT$,

$$\sqrt{n_1}\,\|b_{NT}\| = O_p(\sqrt{T/N}) = o_p(1) \tag{C.28}$$

under the stated aspect-ratio condition.

Consistency of all blocks follows from the exact identities (C.7), (C.9), (C.16), and (C.22), together with the bounds in (C.8), (C.19), and (C.24). These results require only Assumptions 1, 2, and 3. The eigenvalue gap in Assumption 2(iii) yields the required perturbation bounds. Assumption 4 is used only to obtain the centered Gaussian limits. Applying its joint influence-array central limit theorem to (C.10), (C.17), (C.20), and (C.27) gives (24). Multiplying the exposure and loading blocks by H^{-1} gives the coordinates used in $\widehat{\vartheta}$.

With $M \geq 1$ macro types the identity that carries the rate result must be restated for the stacked calendar. Collect the predetermined macro calendar in $\mathbb{D} := (\boldsymbol{D}_0^m, \ldots, \boldsymbol{D}_{T-1}^m)' \in \mathbb{R}^{T\times M}$ and write $\boldsymbol{\delta} := (\delta_1, \ldots, \delta_M) \in \mathbb{R}^{L\times M}$. Stage 2 is the multivariate least-squares fit $\widehat{\boldsymbol{\delta}} = \widehat{A}M\mathbb{D}(\mathbb{D}'M\mathbb{D})^{-1}$, with $\mathbb{D}'M\mathbb{D}$ nonsingular by the stacked calendar second-moment condition of Assumption 3(iii), and the Stage-3 input is $\widehat{A}^0 = \widehat{A} - \widehat{\boldsymbol{\delta}}\,\mathbb{D}'$. Its normal equations are the $L \times M$ identity

$$\widehat{A}^0 M\mathbb{D} = \widehat{A}M\mathbb{D} - \widehat{\boldsymbol{\delta}}(\mathbb{D}'M\mathbb{D}) = 0,$$

so with $\Delta^{(2)} = (\widehat{\boldsymbol{\delta}} - \boldsymbol{\delta})\mathbb{D}'$ we get $\Delta^{(2)}M\widehat{F} = (\widehat{\boldsymbol{\delta}} - \boldsymbol{\delta})\mathbb{D}'M\widehat{A}^{0\prime}\widehat{B} = 0$ and (C.14) holds verbatim. Every bound above then carries over with $M_{[\mathbf{1}_T,\mathbb{D}]}$ in place of $M_{[\mathbf{1}_T,d]}$ in (C.15), the perturbation being still $O_p(T^{-1/2})$, still supported in $\mathrm{col}(B)$, and still absorbed by H.

For a fixed number of firm-event types, replace X_{it}^H in (C.22)–(C.27) with the stacked design interacted with the event-type indicators. The same $\widetilde{\epsilon}_t$ and factor estimation error enter each active block. The covariance matrix can therefore have nonzero off-diagonal blocks even when two event types never occur on the same firm-day. Nonsingularity of the stacked Gram matrix and the fixed

number of event types ensure that the remainder bounds continue to hold. □

C.3 Proof of Theorem 2

Proof. Let $\widehat{U}_{it}^{(0)} = R_{it} - Z'_{i,t-1}\widehat{a}_t$ on the control sample. Expansion of the weighted Stage-1 normal equations gives

$$\widehat{a}_t^\star - \widehat{a}_t = \widehat{Q}_{0t}^{-1} \frac{1}{N_t^{(0)}} \sum_{i:D_{i,t-1}=0} (w_i^\star - 1) Z_{i,t-1} \widehat{U}_{it}^{(0)} + \rho_{1t}^\star. \tag{C.29}$$

The inverse-Gram identity, the maximum-influence condition, and square consistency of the feasible control score make the aggregate contribution of $\rho_{1t}^\star$ $o_p^\star(1)$ at every normalization in $\mathcal{R}_{NT}$.

Stage 2 is a smooth OLS map with denominator bounded away from zero. Its expansion is the projection of (C.29) plus

$$\widehat{v}_m^{-1} \frac{1}{T} \sum_t \widetilde{d}_t w_t^\star \widehat{B}\widehat{B}' \widehat{v}_t, \tag{C.30}$$

and a conditionally negligible remainder. The summands of (C.30) are conditionally independent across t with conditional variance

$$\frac{1}{\widehat{v}_m^2 T^2} \sum_t \widetilde{d}_t^2 \, \widehat{B}\widehat{B}' \widehat{v}_t \widehat{v}_t' \widehat{B}\widehat{B}',$$

and $\widehat{v}_t = B(\bar{F}_t - \mu) + \widetilde{\epsilon}_t + o_p(1)$ together with $\widehat{B}\widehat{B}' \to_p BB'$ makes T times this matrix converge in probability, up to an $O_p(T^{-1} + N^{-1})$ term, to

$$V_\delta := v_m^{-2} \lim_{T\to\infty} \frac{1}{T} \sum_t \mathbb{E}\big[(d_t - \pi_m)^2 \, B(\bar{F}_t - \mu)(\bar{F}_t - \mu)' B'\big],$$

the demeaned weight being read at $\bar{D}_m = \pi_m$ so that the expectation involves only adapted quantities. Here V_δ is the δ block of $\mathcal{V}$ implied by (C.10). The conditional Lindeberg central limit theorem for bounded-moment multipliers (Mammen, 1993; CRW, Lemma C.5) applies to this sum.

For Stage 3, let $\mathcal{B}(S)$ denote the ordered eigenspace map with the regression alignment used in (C.13). On the eigenvalue-gap neighborhood,

$$\widehat{B}^\star - \widehat{B}H^\star = D\mathcal{B}_{\widehat{S}}[\widehat{S}^\star - \widehat{S}] + O_p^\star(\|\widehat{S}^\star - \widehat{S}\|^2), \tag{C.31}$$

whose linear term is the firm-multiplier version of (C.17).

The Stage-2 normal equations annihilate the calendar direction exactly in the draw. Write $\widehat{A}^\star = (\widehat{a}_1^\star, \ldots, \widehat{a}_T^\star)$ and let

$$\widehat{\delta}^{\star,N} := \widehat{v}_m^{-1} \frac{1}{T} \sum_t \widetilde{d}_t (\widehat{a}_t^\star - \bar{\widehat{a}}^\star) = \widehat{A}^\star M d \, (d' M d)^{-1}$$

be the Stage-2 coefficient obtained from the firm-weighted slopes alone. The Stage-3 input is

$\widehat{A}^{0,\star} = \widehat{A}^{\star} - \widehat{\delta}^{\star,N} d'$, so

$$\widehat{A}^{0,\star} M d = \widehat{A}^{\star} M d - \widehat{\delta}^{\star,N}(d' M d) = 0 \tag{C.32}$$

by the definition of $\widehat{\delta}^{\star,N}$, exactly as $\widehat{A}^{0} M d = 0$ holds in the sample. Writing $\Delta^{(2)\star} := (\widehat{\delta}^{\star,N} - \widehat{\delta}) d'$ for the residualization difference and $\widehat{F}^{\star} = \widehat{A}^{0,\star\prime}\widehat{B}^{\star}$ for the draw's factor path, it follows that $\Delta^{(2)\star} M \widehat{F}^{\star} = (\widehat{\delta}^{\star,N} - \widehat{\delta})\, d' M \widehat{A}^{0,\star\prime}\widehat{B}^{\star} = 0$, so the bootstrap analogue of (C.14) is exact. The date multiplier generates the Stage-2 draw in (C.30) and reaches Stage 3 only through the macro residualization, which is carried out with $\widehat{\delta}^{\star,N}$.

Applying the same derivative to $(I_L - \widehat{B}\widehat{B}')\bar{\widehat{a}}^{0}$ and to $\widehat{B}'\widehat{a}_t^{0}$ reproduces (C.20) and (C.21). The conditional quadratic eigenvector and generated-design terms are $O_p^{\star}(N^{-1})$, so after $\sqrt{NT}$ or $\sqrt{n_1}$ normalization they are $O_p^{\star}(\sqrt{T/N}) = o_p^{\star}(1)$.

A Taylor expansion of the firm-weighted Stage-4a normal equations around $(\widehat{\theta}_{\gamma}, \widehat{a}^{\star}, \widehat{F})$, with $\widehat{U}_{it} := R_{it} - Z_{i,t-1}'\widehat{a}_t - \widehat{X}_{it}'\widehat{\theta}_{\gamma}$ the Stage-4a fitted residual on $\mathcal{T}^{(1,0)}$, gives the bootstrap counterparts of the three scores in (C.27),

$$\begin{aligned}
&\frac{1}{\sqrt{n_1}} \sum_{(i,t)\in\mathcal{T}^{(1,0)}} (w_i^{\star} - 1)\widehat{X}_{it}\widehat{U}_{it}, \\
&- \frac{1}{\sqrt{n_1}} \sum_{t} \widehat{W}_t(\widehat{a}_t^{\star} - \widehat{a}_t), \\
&- \frac{1}{\sqrt{n_1}} \sum_{t} \widehat{A}_t\{\widehat{F}_t^{\star} - (H^{\star})^{-1}\widehat{F}_t\}.
\end{aligned} \tag{C.33}$$

Because the same firm-weighted Stage-1 draw enters the response adjustment and the re-estimated PCA, (C.33) reproduces the Stage-1 and Stage-3 cross covariance. Centering the loading draw at $\widehat{B}H^{\star}$ and the exposure draw at $\widehat{\Gamma}H^{\star}$ removes only the coordinate change and not the loading-space error, because the regression alignment $H^{\star}$ plays for the pair $(\widehat{B}, \widehat{B}^{\star})$ exactly the role that H in (C.13) plays for $(B, \widehat{B})$. No first-order block of the target law is annihilated.

Collecting (C.29) to (C.33), including the stacked blocks when present, gives feasible influence arrays $\widehat{\phi}_{t,NT}$ and $\widehat{\psi}_{i,NT}$ such that

$$\mathcal{R}_{NT}(\widehat{\vartheta}^{\star} - \widehat{\vartheta}) = \sum_{t=1}^{T} w_t^{\star}\widehat{\phi}_{t,NT} + \sum_{i=1}^{N}(w_i^{\star} - 1)\widehat{\psi}_{i,NT} + o_p^{\star}(1). \tag{C.34}$$

The two multiplier sums are conditionally independent, and given the sample the summands $(w_i^{\star} - 1)\widehat{\psi}_{i,NT}$ are independent across i and exactly centered. For every fixed conformable vector ν,

$$\mathrm{Var}^{\star}\left(\sum_{t} w_t^{\star}\nu'\widehat{\phi}_{t,NT} + \sum_{i}(w_i^{\star} - 1)\nu'\widehat{\psi}_{i,NT}\right) \to_p \nu'\mathcal{V}\nu. \tag{C.35}$$

Indeed, the left side is the sum of the two conditional variances, because the multiplier arrays are independent, so it can converge to $\nu'\mathcal{V}\nu$ only if the population cross covariance of the date and firm

influence sums vanishes, which is Assumption 5(iii). Given that condition, the conditional variance is $\sum_t(\nu'\widehat{\phi}_{t,NT})^2+\sum_i(\nu'\widehat{\psi}_{i,NT})^2$. Square consistency transfers each sum to its infeasible counterpart by Cauchy–Schwarz, using the $O_p(1)$ bound on the infeasible sums supplied by the outer-product convergence of Assumption 5(ii), and those infeasible sums converge to $\nu'\mathcal{V}\nu$ by the same limit.

The maximum-influence and multiplier-moment conditions give the conditional Lindeberg condition. Conditional Lindeberg–Feller and Cramér–Wold therefore imply that (C.34) converges, in probability in the bounded-Lipschitz metric, to $\mathcal{N}(0,\mathcal{V})$, the same law as $\mathcal{R}_{NT}(\widehat{\vartheta}-\vartheta)$. The matrix $\mathcal{V}$ is singular, since by (C.10) the $\sqrt{T}$-scaled δ block has a limit supported on $\mathrm{col}(B)$, of rank K, and the α and loading blocks jointly satisfy the K restrictions (27). The limiting law is therefore concentrated on a proper subspace of $\mathbb{R}^p$ and has no Lebesgue density.

For a fixed number of event types the same calculation applies to the stacked arrays, and the Stage-2 annihilation is derived as in the sample. With $\widehat{\boldsymbol{\delta}}^{\star,N} := \widehat{A}^{\star}M\mathbb{D}(\mathbb{D}'M\mathbb{D})^{-1}$,

$$\widehat{A}^{0,\star}M\mathbb{D} = \widehat{A}^{\star}M\mathbb{D} - \widehat{\boldsymbol{\delta}}^{\star,N}(\mathbb{D}'M\mathbb{D}) = 0,$$

so (C.32) holds with the stacked calendar in place of d. □

C.4 Proof of Corollary 1

Proof. Put $H_{NT} := \mathcal{R}_{NT}(\widehat{\vartheta}-\vartheta)$, $H^{\star}_{NT} := \mathcal{R}_{NT}(\widehat{\vartheta}^{\star}-\widehat{\vartheta})$, and $\nu_{NT} := \sigma^{-1}_{NT}\mathcal{R}^{-1}_{NT}\nabla g(\vartheta)$, so that $\nu'_{NT}\mathcal{V}\nu_{NT}=1$. Since $H_{NT}=O_p(1)$ by Theorem 1, condition (c) gives

$$g(\widehat{\vartheta})-g(\vartheta) = \sigma_{NT}\nu'_{NT}H_{NT}+o_p(\sigma_{NT}). \tag{C.36}$$

We first deal with the singularity of $\mathcal{V}$. Let $P_{\mathcal{V}}$ denote the orthogonal projector onto range$(\mathcal{V})$. Since $\nu'_{NT}\mathcal{V}\nu_{NT} = (P_{\mathcal{V}}\nu_{NT})'\mathcal{V}(P_{\mathcal{V}}\nu_{NT}) = 1$, and the positive eigenvalues of $\mathcal{V}$ are bounded away from zero, $P_{\mathcal{V}}\nu_{NT}$ is bounded. The null space consists of the $\ker(B')$ directions of the δ block and the K joint (α,B) directions of (27). For the latter, (27) is the first-order implication of the exact sample restriction $\widehat{\alpha}'\widehat{B}=0$, so the corresponding projections of H_{NT} are $o_p(1)$. Their coefficients in ν_{NT} are bounded because these blocks have rate $\sqrt{NT}$, $n_1 \asymp NT$, and a positive first-order standard deviation cannot have order faster than $(NT)^{-1/2}$.

For the δ null directions, $(I_L-BB')\sqrt{T}(\widehat{\delta}-\delta)=o_p(1)$ by Theorem 1. If $\nabla_{\delta}g(\vartheta)\in\mathrm{col}(B)$, their contribution vanishes. Otherwise condition (b) implies that σ_{NT} is of order $T^{-1/2}$, so $\frac{\|(I_L-BB')\nabla_{\delta}g(\vartheta)\|}{\sigma_{NT}\sqrt{T}} = O(1)$, and the contribution is again $o_p(1)$. Consequently, $\nu'_{NT}H_{NT} = (P_{\mathcal{V}}\nu_{NT})'H_{NT}+o_p(1)$. Since $P_{\mathcal{V}}\nu_{NT}$ is bounded and has variance normalized to one, Theorem 1 gives

$$\nu'_{NT}H_{NT} \Rightarrow \mathcal{N}(0,1),$$

where boundedness permits the usual subsequence argument if $P_{\mathcal{V}}\nu_{NT}$ has no limit. Together with

(C.36),

$$\sigma_{NT}^{-1}\{g(\widehat{\vartheta}) - g(\vartheta)\} \Rightarrow \mathcal{N}(0,1). \tag{C.37}$$

For the bootstrap, Theorem 2 gives $H_{NT}^{\star} = O_p^{\star}(1)$ conditionally on the sample and in probability. Define

$$r_{NT}(h) := g(\vartheta + \mathcal{R}_{NT}^{-1}h) - g(\vartheta) - \nabla g(\vartheta)'\mathcal{R}_{NT}^{-1}h.$$

Condition (c) and the tightness of H_{NT} and $H_{NT}^{\star}$ imply $r_{NT}(H_{NT}) = o_p(\sigma_{NT})$, $r_{NT}(H_{NT}+H_{NT}^{\star}) = o_p^{\star}(\sigma_{NT})$ conditionally and in probability. Hence

$$\begin{aligned} g(\widehat{\vartheta}^{\star}) - g(\widehat{\vartheta}) &= \sigma_{NT}\nu_{NT}'H_{NT}^{\star} + r_{NT}(H_{NT} + H_{NT}^{\star}) - r_{NT}(H_{NT}) \\ &= \sigma_{NT}\nu_{NT}'H_{NT}^{\star} + o_p^{\star}(\sigma_{NT}). \end{aligned} \tag{C.38}$$

The same null-space argument applies to $H_{NT}^{\star}$ as the joint (α, B) null projections are $o_p^{\star}(1)$, and condition (b) removes any first-order contribution from the $\ker(B')$ directions when σ_{NT} is faster than $T^{-1/2}$. The conditional bounded-Lipschitz convergence of Theorem 2 therefore gives $\sigma_{NT}^{-1}\{g(\widehat{\vartheta}^{\star}) - g(\widehat{\vartheta})\} \Rightarrow^{\star} \mathcal{N}(0,1)$ conditionally on the sample and in probability.

Both scalar laws thus converge to the same continuous $\mathcal{N}(0,1)$ distribution, so their distribution functions converge uniformly and the difference in the statement vanishes. Percentile and studentized intervals follow. Finally, the bounded $(4+\varepsilon)$ moments of Assumptions 3(ii) and 5(i), together with the expansion above, give uniform integrability of the squared normalized bootstrap statistic. Its conditional second moment therefore converges to one, and the bootstrap variance of $g(\widehat{\vartheta})$ is consistent for σ_{NT}^2. □

C.5 Proof of Corollary 2

Proof. Averaging the Stage-1 identity (C.7) over dates and using the Stage-2 definitions,

$$\bar{\widehat{a}}^0 = \alpha + B\mu_T + \bar{\widetilde{\epsilon}} - (\widehat{\boldsymbol{\delta}} - \boldsymbol{\delta})\,\bar{\mathbb{D}}, \tag{C.39}$$

where $\bar{\mathbb{D}} \in \mathbb{R}^M$ collects the sample means of the macro calendars. Let $\Xi_T := (\mathbb{D}'M\mathbb{D})^{-1}\mathbb{D}'M\bar{F} \in \mathbb{R}^{M\times K}$ be the calendar coefficient of the factor path, so that $\mu_T^0 = \mu_T - \Xi_T'\bar{\mathbb{D}}$ by definition. Substituting $\widehat{A} = A + \widetilde{\mathcal{E}}$ with $\widetilde{\mathcal{E}} := (\widetilde{\epsilon}_1, \ldots, \widetilde{\epsilon}_T)$ and $A = \alpha\mathbf{1}_T' + B\bar{F}' + \boldsymbol{\delta}\,\mathbb{D}'$ into the Stage-2 formula gives the exact decomposition

$$\widehat{\boldsymbol{\delta}} - \boldsymbol{\delta} = B\,\Xi_T' + \mathbb{E}, \qquad \mathbb{E} := \widetilde{\mathcal{E}}M\mathbb{D}(\mathbb{D}'M\mathbb{D})^{-1}, \tag{C.40}$$

because $\alpha\mathbf{1}_T'M = 0$. Combining (C.39) and (C.40) gives $\bar{\widehat{a}}^0 = \alpha + B\mu_T^0 + \bar{\widetilde{\epsilon}} - \mathbb{E}\,\bar{\mathbb{D}}$ exactly. The in-span calendar coefficient of the factor path is absorbed into the baseline mean μ_T^0, and only the calendar coefficient of the Stage-1 noise remains. With $\widehat{B} = BH + W_B$, $B'B = I_K$ and $B'\alpha = 0$,

$$\widehat{\mu} - H^{-1}\mu_T^0 = (H' - H^{-1})\mu_T^0 + W_B'(\alpha + B\mu_T^0) + \widehat{B}'(\bar{\widetilde{\epsilon}} - \mathbb{E}\,\bar{\mathbb{D}}), \tag{C.41}$$

again exactly. The first term equals $-sH^{-1}\mu_T^0$ with $s := H'B'W_B + W_B'BH + W_B'W_B$, which is linear in W_B up to $O_p(\|W_B\|^2)$, and by (C.17) $\sqrt{NT}\,W_B$ is a fixed linear image of S_B plus a bias of order $\sqrt{T/N}$ that is negligible under $T/N \to 0$. Next, $\sqrt{NT}\,\bar{\widetilde{\epsilon}} = S_\epsilon$ exactly, and $\sqrt{NT}\,\mathbb{E}$ is a fixed linear image of the date array $\{\widetilde{d}_t\sqrt{N}\,\widetilde{\epsilon}_t\}$, up to replacing the sample calendar Gram $T^{-1}\mathbb{D}'M\mathbb{D}$ by its probability limit, which is the convention already used in forming (26). Its orthogonal projection is $S_\delta^\perp$, and its in-span component is covered by condition (a). Hence $\sqrt{NT}(\widehat{\mu} - H^{-1}\mu_T^0)$ converges jointly with the blocks of (24) as a fixed bounded linear image of the same Gaussian limit.

The bilinear decomposition

$$\begin{aligned}\widehat{c}_z - c_z = {} & z'(\widehat{\gamma}_j - \gamma_j) + z'(\widehat{\Gamma}_j - \Gamma_j H)H^{-1}\mu_T^0 \\ & + z'\Gamma_j H(\widehat{\mu} - H^{-1}\mu_T^0) + z'(\widehat{\Gamma}_j - \Gamma_j H)(\widehat{\mu} - H^{-1}\mu_T^0) \\ & + z'\Gamma_j(\mu_T^0 - \mu)\end{aligned} \tag{C.42}$$

holds exactly. The first three terms are $O_p(n_1^{(j)-1/2})$, $O_p(n_1^{(j)-1/2})$ and $O_p((NT)^{-1/2} + N^{-1})$. The fourth is the product of two of these rates, and since $n_1^{(j)} \asymp NT$ by Assumption 3(iii) it is $o_p(\sigma_{c,NT})$ whenever $\sigma_c > 0$. The path term is $O_p(T^{-1/2})$. For that term write $\mu_T^0 - \mu = (\mu_T - \mu) - \Xi_T'\bar{\mathbb{D}}$. Under Assumption 1(i), $\sqrt{T}(\mu_T - \mu)$ is a martingale-difference average of $\{\bar{F}_t - \mu\}$ and $\sqrt{T}\,\Xi_T$ is a fixed linear image of the date array $\{\widetilde{d}_t(\bar{F}_t - \mu)\}$, so $\sqrt{T}(\mu_T^0 - \mu)$ is a fixed bounded linear image of the date-level path score that condition (a) includes, jointly with the blocks of (24) and with the calendar-noise score. Applying the joint central limit theorem to the four first-order terms, namely the (γ, Γ) blocks of (C.27), the expansion (C.41) and the path score, yields $\sigma_{c,NT}^{-1}(\widehat{c}_z - c_z) \Rightarrow \mathcal{N}(0,1)$, the normalization $n_1^{(j)} \wedge T$ of condition (b) keeping every retained term inside the limit and the cross term negligible at either rate.

For the bootstrap, every draw recomputes $\widehat{\mu}^\star = T^{-1}\sum_t \widehat{F}_t^\star = \widehat{B}^{\star\prime}\bar{\widehat{a}}^{0,\star}$ from its own stages, and the conditional analogue of (C.39) to (C.41) holds with the firm-weighted arrays. Because the Stage-3 input is residualized with the firm-only draw $\widehat{\boldsymbol{\delta}}^{\star,N}$, the draw's calendar-noise term is $(\widehat{\boldsymbol{\delta}}^{\star,N} - \widehat{\boldsymbol{\delta}})\bar{\mathbb{D}}$, whose conditional law reproduces that of $\mathbb{E}\,\bar{\mathbb{D}}$. The factor-path component of the calendar coefficient is held fixed by the draw and is redrawn instead through the path-mean term $W^\star$. For that term, conditionally on the sample $\mathbb{E}^\star[\sqrt{T}\,W^\star] = 0$ and

$$\mathrm{Var}^\star(\sqrt{T}\,W^\star) = \frac{T}{T_0^2}\sum_{t:\,D_{m,t-1}=0}(\widehat{F}_t^\star - \bar{\widehat{F}}_0^\star)(\widehat{F}_t^\star - \bar{\widehat{F}}_0^\star)' \xrightarrow{p} \mathrm{Var}(\sqrt{T}(\mu_T^0 - \mu)),$$

the Lindeberg condition holding by the bounded fourth moments of Assumption 3(ii). Because $W^\star$ and the Stage-2 calendar draw share the multiplier $w_t^\star$, their conditional covariance is the sample cross-moment of the corresponding date arrays, which converges to the population covariance of condition (a). The remaining blocks are those of (C.33) and (C.34), so the conditional Lindeberg argument in the proof of Theorem 2 applies to the augmented linear image. The cross term in (C.42) is $o_p^\star(\sigma_{c,NT})$ by the same rate count, and the delta-method step in the proof of Corollary 1 transfers the scalar law. Percentile and studentized validity follow as there. □

C.6 Proof of Proposition 2

Proof. *(i)* Collect the conclusion of the proof of Theorem 1 as the unscaled expansion

$$\widehat{\vartheta} - \vartheta = \Lambda_{N,T} + \mathbf{b}_N + \mathbf{r}_{N,T}, \tag{C.43}$$

where $\Lambda_{N,T}$ is the influence-linear statistic, each block a fixed population-normalized linear image of the score arrays of (25) with the population Grams Q_{0t} and Σ_4. The term $\mathbf{b}_N$ stacks the bias terms of (C.17), (C.20) and (C.28), and $\mathbf{r}_{N,T}$ collects the stochastic remainders.

The remainder bounds in the proof of Theorem 1 use $T/N \to 0$ nowhere except through $\mathbf{b}_N$. The $o_p(N^{-1})$ statements displayed there are the $T/N \to 0$ specializations of finer orders that their derivations deliver. The centered quadratic terms have conditional standard deviation $O((N^2T)^{-1/2})$, the cross terms are $O_p(T^{-1/2}N^{-1})$, and the products of type $O_p(\|W_B\|^2)$ and $O_p(T^{-1/2}\|W_B\|)$ are $O_p(N^{-2})$ up to factors that vanish after scaling. In those forms every remainder is o_p of its block's rate whenever $T = o(N^3)$, which condition (b) implies.

Each half cross-section satisfies the same conditions with $N/2$ in place of N. The realized half Grams are nonsingular with probability approaching one by Assumption 3(i), the sampling shares halve proportionally, and the alignment of the half-sample basis to the full-sample basis composes with H, which the aligned blocks do not feel. Fix a partition s and write $\widehat{\vartheta}_{(h)} := \widehat{\vartheta}_{(s,h)}$. By condition (a),

$$\widehat{\vartheta}_{(h)} - \vartheta = \Lambda^{(h)}_{N,T} + 2\,\mathbf{b}_N + O_p(N^{-3/2}) + \mathbf{r}^{(h)}_{N,T}, \qquad h = 1, 2.$$

The influence-linear statistics recombine exactly. For any block whose score is a normalized sum over firms or firm-dates, which covers the managed-portfolio, loading-space, orthogonal-calendar and Stage-4a components, the population normalization is common to the full panel and to the halves, so $\frac{1}{2}(\Lambda^{(1)}+\Lambda^{(2)}) = \Lambda_{N,T}$ up to the replacement of half by full cell counts. That replacement is an $O_p(N^{-1/2})$-relative perturbation, because the firms are split at random and the half cell counts satisfy $n_1^{(h)} = (n_1/2)(1 + O_p(N^{-1/2}))$, so after $\sqrt{NT}$ scaling it contributes $O_p(N^{-1/2}) = o_p(1)$. The calendar score is common to both halves and recombines identically. Hence the per-partition combination $\widehat{\vartheta}^{\mathrm{jk}}_{(s)} := 2\,\widehat{\vartheta} - \frac{1}{2}(\widehat{\vartheta}_{(s,1)} + \widehat{\vartheta}_{(s,2)})$ satisfies

$$\widehat{\vartheta}^{\mathrm{jk}}_{(s)} - \vartheta = \Lambda_{N,T} + O_p(N^{-3/2}) + \widetilde{\mathbf{r}}_{N,T},$$

with $\widetilde{\mathbf{r}}_{N,T}$ obeying the same bounds as $\mathbf{r}_{N,T}$. After scaling, the bias residual is $O_p(\sqrt{NT}\,N^{-3/2}) = O_p(\sqrt{T/N^2}) = o_p(1)$, and the joint central limit theorem of Assumption 4(ii), which references no aspect ratio, applied to $\Lambda_{N,T}$ gives the stated limit with the covariance $\mathcal{V}$ of (24). Since $\widehat{\vartheta}^{\mathrm{jk}} = S^{-1}\sum_s \widehat{\vartheta}^{\mathrm{jk}}_{(s)}$, averaging over the S independent partitions doesnt change the results. The influence recombination is exact partition by partition, the bias cancellation holds per partition with the common coefficient b, and the second-order remainders of distinct partitions obey the same bounds.

(ii) The single use of $T/N \to 0$ on the draw side of the proof of Theorem 2 is the bound on

the conditional quadratic eigenvector and generated-design terms. Decompose each such term into its conditional mean given the sample and a multiplier-centered fluctuation. The conditional mean reproduces the sample's own bias coefficient $\widehat{\mathbf{b}}_N$ and is common to every draw, so it drops out of the mean-corrected bootstrap covariance exactly and does not affect intervals centered at $\widehat{\vartheta}^{\mathrm{jk}}$. The centered fluctuation is a degenerate quadratic form in the firm multipliers whose coefficient structure splits in two. The N diagonal coefficients are of order N^{-2} and contribute conditional variance $O_p(N^{-3})$ by the multiplier fourth moments of Assumption 5(i). The off-diagonal coefficients are time-averaged cross products of conditionally uncorrelated firm arrays, hence of order $N^{-2}T^{-1/2}$ by the conditional cross-firm independence and moment bounds of Assumption 3(ii), and contribute conditional variance $O_p(N^{-2}T^{-1})$. The conditional standard deviation is therefore $O_p^\star(N^{-3/2} + N^{-1}T^{-1/2})$ blockwise, contributing $O_p^\star(\sqrt{T/N^2}) + O_p^\star(N^{-1/2}) = o_p^\star(1)$ after scaling under condition (b). The conditional Lindeberg argument of Step 2 of the proof of Theorem 2 then delivers the same Gaussian limit $\mathcal{N}(0, \mathcal{V})$ conditionally, so the full-sample bootstrap standard errors remain consistent for the jackknifed blocks, and coverage of intervals centered at $\widehat{\vartheta}^{\mathrm{jk}}$ follows from the two limits as in the proof of Corollary 1. □

C.7 Proof of Proposition 3

Proof. (i) Write $W_B := \hat{B} - BH$, $P := B'\hat{B}$ and $\Psi := I_K - HP'$ as in the proof of Theorem 1. On a coincidence date the Stage-3 input is $\hat{a}_t^0 = \alpha + B\bar{F}_t + \widetilde{\epsilon}_t - \sum_g (\hat{\delta}_g - \delta_g) D_{m,t-1}^{(g)}$, so $\hat{F}_t = \hat{B}'\hat{a}_t^0$, $B'\alpha = 0$ and $B'B = I_K$ give the exact identity

$$\hat{F}_t = H^{-1}\bar{F}_t - \sum_{g=1}^{M} H'B'(\hat{\delta}_g - \delta_g)\, D_{m,t-1}^{(g)} + e_t, \qquad e_t := \hat{B}'\widetilde{\epsilon}_t + W_B'\alpha - H^{-1}\Psi\bar{F}_t - \sum_{g=1}^{M} W_B'(\hat{\delta}_g - \delta_g) D_{m,t-1}^{(g)},$$

the first three terms of e_t being those of (C.21) and the fourth $O_p(\|W_B\|\, T^{-1/2})$. Inverting,

$$\bar{F}_t = H\hat{F}_t + \sum_{g=1}^{M} HH'B'(\hat{\delta}_g - \delta_g)\, D_{m,t-1}^{(g)} - He_t.$$

Substituting this into the coincidence-cell conditional mean of (14), and using $\tilde{R}_{it} = q_{it} - Z_{i,t-1}'\widetilde{\epsilon}_t$,

$$\tilde{R}_{it} = \sum_{g=1}^{M} Z_{i,t-1}'\{c_g + \Gamma_a HH'B'(\hat{\delta}_g - \delta_g)\} D_{m,t-1}^{(g)} + Z_{i,t-1}'(\Gamma_a H)\hat{F}_t + \varepsilon_{it},$$

with $\varepsilon_{it} := U_{it}(1) - Z_{i,t-1}'\widetilde{\epsilon}_t - Z_{i,t-1}'\Gamma_a He_t$. The Stage-4b level regressors already span the calendar translation. For macro type g, it is absorbed by the coefficient on $D_{m,t-1}^{(g)} Z_{i,t-1}$ and does not affect the exposure coefficient on $Z_{i,t-1} \otimes \hat{F}_t$. When several releases occur on the same date, the corresponding level shifts enter additively, as specified in Section 2.4. Solving the Stage-4b normal

equations under conditions (a) and (b), and using $HH' = I_K + O_p(\|W_B\|)$,

$$\hat{c}_g - c_g = \Gamma_a B'(\hat{\delta}_g - \delta_g) + r_{c,g}, \qquad \check{\Gamma}_a - \Gamma_a = (\hat{\Gamma}_a - \Gamma_a H)H^{-1},$$

where $r_{c,g}$ collects the Stage-4b score in ε_{it}, the generated-design bias, and $\Gamma_a(HH' - I_K)B'(\hat{\delta}_g - \delta_g) = O_p(\|W_B\|\, T^{-1/2})$. For the minimum-distance block, $\check{\lambda}_g = H\hat{B}'\hat{\delta}_g$ and $\hat{B}' = H'B' + W_B'$ give

$$\check{\lambda}_g - \lambda_g = B'(\hat{\delta}_g - \delta_g) + r_{\lambda,g}, \qquad r_{\lambda,g} = (HH' - I_K)\lambda_g + HW_B'\hat{\delta}_g + O_p(\|W_B\|\, T^{-1/2}).$$

By (C.19), $\|W_B\| = O_p((NT)^{-1/2} + N^{-1})$, so under $T/N \to 0$ every term of $r_{c,g}$, $r_{\lambda,g}$ and $\check{\Gamma}_a - \Gamma_a$ is $O_p(n_{1m}^{-1/2})$, with the generated-design biases negligible after $\sqrt{n_{1m}}$ scaling exactly as in the proof of Theorem 1. Condition (c) gives their joint Gaussian limit, and condition (a) makes $n_{1m}^{(g)} \asymp n_{1m}$, so the common scaling applies to every reported type.

(ii) Since $\check{\Gamma}_a\check{\lambda}_g = \hat{\Gamma}_a\hat{\lambda}_g$,

$$\hat{\gamma}_{a,g} - \gamma_{a,g} = (\hat{c}_g - c_g) - \Gamma_a(\check{\lambda}_g - \lambda_g) - (\check{\Gamma}_a - \Gamma_a)\lambda_g - (\check{\Gamma}_a - \Gamma_a)(\check{\lambda}_g - \lambda_g).$$

By (i) the first two terms are $\Gamma_a B'(\hat{\delta}_g - \delta_g) + r_{c,g} - \Gamma_a\{B'(\hat{\delta}_g - \delta_g) + r_{\lambda,g}\} = r_{c,g} - \Gamma_a r_{\lambda,g}$, so the $O_p(T^{-1/2})$ translation cancels exactly, and the last term is $O_p(n_{1m}^{-1/2}) \times O_p(T^{-1/2}) = o_p(n_{1m}^{-1/2})$. The expansion follows, and joint asymptotic normality of $\sqrt{n_{1m}}(\hat{\gamma}_{a,g} - \gamma_{a,g}, \mathrm{vec}(\check{\Gamma}_a - \Gamma_a))$, jointly over the reported types, is the corresponding linear image of the limit in (i).

(iii) Stage 3 is residualized with the firm-only coefficient $\hat{\delta}^{\star,N}$, so the draw's factor path satisfies the analogue of the identity of (i),

$$\hat{F}_t^\star = (H^\star)^{-1}\hat{F}_t - \sum_{g=1}^{M} H^{\star\prime}\hat{B}'(\hat{\delta}_g^{\star,N} - \hat{\delta}_g)\, D_{m,t-1}^{(g)} + e_t^\star,$$

with $e_t^\star$ the multiplier analogue of e_t. The translation is again collinear with the level regressors, so it enters $\hat{c}_g^\star$ with coefficient $\hat{\Gamma}_a$; and because $\hat{\lambda}_g^{\star,N}$ is formed from the same $\hat{\delta}_g^{\star,N}$, it carries the translation with coefficient I_K. The two cancel identically in $\hat{\gamma}_{a,g}^\star$, exactly as in (ii). By the construction of Section B.3, the reported draw $\hat{\delta}_g^\star$ carries the date multiplier, so $\hat{\delta}_g^\star - \hat{\delta}_g^{\star,N} = O_p^\star(T^{-1/2})$ while the Stage-3 input does not, and the difference would survive the combination as an $O_p^\star(\sqrt{N})$ term after $\sqrt{n_{1m}}$ scaling. Square consistency of the feasible split-weighted Stages-1–4b influence array, condition (c), then delivers conditional convergence of the centered pair to the limit of (ii), by the conditional Lindeberg argument of Step 2 of the proof of Theorem 2.

For the announcement-day effect, let μ_T^0 and Ξ_T be the intercept and slopes of the least-squares fit of $\bar{F}_t$ on x_t, as in the proof of Corollary 2. By (C.40), $B'(\hat{\delta}_g - \delta_g) = \Xi_{g,T} + B'\mathbb{E}_g$, and by (C.41), $\hat{\mu} = H^{-1}\mu_T^0 + O_p((NT)^{-1/2} + N^{-1})$, so (i) gives

$$\hat{c}_g + \hat{\Gamma}_a\hat{\mu} - (c_g + \Gamma_a\mu) = \Gamma_a(\mu_T^0 + \Xi_{g,T} - \mu) + O_p(n_{1m}^{-1/2}).$$

The $\sqrt{T}$ component is thus the type-g calendar level of the path: for $M = 1$ it is $\Gamma_a(\bar{F}_1 - \mu)$ with $\bar{F}_1$ the announcement-date average, because $\hat{\mu}$ recovers the clean-date average while $\hat{c}_g$ absorbs the calendar difference. Write $Q_{x,T} := T^{-1}\sum_t x_t x_t'$ and let $\Theta_0 := (\mu, 0, \ldots, 0)'$ be the population coefficient of the regression of $\bar{F}_t$ on x_t, which takes this form because the calendars are predetermined and $\mathbb{E}[\bar{F}_t \mid \mathcal{G}_{t-1}] = \mu$. Since $\bar{F}_t' - x_t'\Theta_0 = (\bar{F}_t - \mu)'$, the sample coefficient $\Theta_T := (\mu_T^0, \Xi_T')'$ satisfies the exact identity

$$\sqrt{T}(\Theta_T - \Theta_0) = Q_{x,T}^{-1}\,\frac{1}{\sqrt{T}}\sum_{t=1}^{T} x_t(\bar{F}_t - \mu)',$$

so with $\ell_g := e_0 + e_g$, $\sqrt{T}(\mu_T^0 + \Xi_{g,T} - \mu) = \ell_g' Q_{x,T}^{-1} T^{-1/2}\sum_t x_t(\bar{F}_t - \mu)'$, a martingale-difference average by Assumption 1(i) because x_t is $\mathcal{G}_{t-1}$-measurable. The draw $W_g^\star$ of Remark 3 is the wild-multiplier version of the same functional. Under the non-overlap condition stated there the projection is saturated, so $W_g^\star = \ell_g' W^\star$ with the fitted residuals $\hat{F}_t^\star - \widehat{\Theta}^{\star\prime} x_t$ as feasible substitutes for $\bar{F}_t - \mu$, and its conditional law reproduces that limit. Sharing $w_t^\star$ with the Stage-2 calendar draw reproduces its covariance with the calendar score. Corollary 2 then applies to $z'(\hat{c}_g + \hat{\Gamma}_a\hat{\mu})$ and its draw. □

D Firm Characteristics

Table 7: Firm characteristics.

Characteristic	Definition	Frequency
Asset growth	One-year growth in total assets	Annual
Profit growth	One-year change in gross profit scaled by total assets	Annual
Investment	Capital expenditures scaled by total assets	Annual
Profit margin	Gross profit scaled by sales	Annual
Return on assets	Gross profit scaled by total assets	Annual
Return on equity	Operating profitability scaled by book equity	Annual
Return on investment	Gross profit scaled by book enterprise value	Annual
Return on physical capital	Gross profit scaled by net property, plant, and equipment	Annual
Issuance	Net financial cash flow scaled by total assets	Annual
Equity payout	Net equity payout scaled by total assets	Annual
Leverage	Total debt scaled by market enterprise value	Monthly
Financial soundness	Interest expense scaled by total debt	Annual
Solvency	Total debt scaled by total assets	Annual
Liquidity	Days inventory outstanding	Annual
Activity	Inventory turnover	Annual
Book-to-market	Book equity scaled by market equity	Monthly
Profit-to-market	Gross profit scaled by market equity	Monthly
Equity payout-to-market	Fiscal-year total dividends scaled by market equity (statement-based; distinct from the trailing measure below)	Monthly
Depreciation growth	One-year depreciation growth scaled by assets	Annual
Size	Market equity	Monthly
Dividend-to-market	Trailing 12-month dividends scaled by market equity	Monthly
Short-term reversal	Return over the most recent completed month	Monthly
Long-term momentum	Cumulative return from months $t-12$ through $t-1$	Monthly
Seasonality	Return in the same calendar month one year earlier	Monthly
Idiosyncratic volatility	Idiosyncratic volatility from a 21-day Fama–French three-factor regression.	Monthly
Return volatility	Daily excess-return volatility over the prior 252 trading days.	Monthly
CAPM beta	CAPM beta estimated from daily returns over the prior 252 trading days.	Monthly
Zero-trade days	Number of zero-volume trading days over the prior 252 trading days.	Monthly
Turnover	Average daily share turnover over the prior 126 trading days.	Monthly

Notes: The table reports the 29 firm characteristics included in the empirical specification. Accounting characteristics are constructed from the most recently available Compustat data, while market characteristics are constructed from CRSP data. Frequency indicates how often each characteristic is updated. Before estimation, each characteristic is rank-normalized by date using average ranks for ties and recentered to have a cross-sectional mean of zero. Section 4.1 describes the variable timing, sample filters, and construction of the final sample.