Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Front-Loaded Innovation, Slow-Moving Level: Climate Policy Uncertainty and the Term Structure of the Green–Brown Return Gap

RunShi Lai

lairunshi@student.zy.cdut.edu.cn

Chengdu University of Technology

Research Article

Keywords: Climate policy uncertainty, Green assets, Asset duration, Local projections, Term structure, China

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11662431/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Climate policy uncertainty determines the cost of capital for the low-carbon transition, yet conventional models mis-measure its transmission to asset prices by entering a single uncertainty index without separating its slow-moving level from its transitory innovation. This paper introduces a text-augmented local-projection framework using a news-based China Climate Policy Uncertainty (CCPU) index and green–brown equity returns. We construct time-aligned panel data from the CCPU index and green–brown portfolio returns, decompose each into a standardized level and first difference, and estimate horizon-by-horizon local projections on one sample and controls; a design-based wild bootstrap replaces the conventional 1.96 with horizon-specific critical values. The innovation component is significant: a one-standard-deviation shock raises the green–brown gap by 2.67 points at impact (t = 2.61, clearing its design-earned threshold of 2.21) and contributes nothing thereafter, while the level component reaches − 16.9 points at h = 12, concentrated on the long-duration green leg (− 19.2 vs − 2.3 on brown) and implies 226 basis points of required return. The wild bootstrap establishes that the five-per-cent threshold is 2.21 for the innovation and 3.97 for the level, so the first is an estimate and the second a description; identifying the persistent and transitory innovation as the distinct drivers of the gap. The framework establishes that a single uncertainty index cannot recover both the impact repricing and the term-structure effect, and that the sign disagreement in the literature is a disagreement about the regressor rather than about markets or periods—so component decomposition is necessary to measure climate-policy risk correctly.

JEL Classification C15 · C22 · C53 · G12 · Q54

1 Introduction

Three computational choices precede any estimate of how a policy innovation propagates, and each is usually settled by convention rather than by argument. A persistent uncertainty index can enter in levels or in first differences, and the two are nearly collinear by construction. Thirteen horizons can be reported as thirteen tests or as one family. And a twelve-month overlapping window on seventy-six observations holds far fewer independent observations than its length suggests. This paper treats the three as objects of study on a case in which they bind, and reports what each one does to the answer.

The case is the pricing of climate policy uncertainty in green relative to brown equity. Green assets are priced against a policy backdrop that changes monthly. The empirical literature on green asset pricing has worked mostly with long panels and cross-sections, asking whether investors accept a lower expected return on green securities (Zerbib, 2019; Pástor et al., 2021) or whether carbon-intensive firms carry a premium (Bolton and Kacperczyk, 2021). Those designs answer a question about the level of expected returns over years. They do not answer what happens in the month when climate policy becomes more uncertain.

The two objects behave differently. A persistent difference in expected returns accumulates, so its imprint on prices widens with the holding horizon. An innovation in policy uncertainty is resolved, so its

effect on relative prices is largest when the news arrives and does not build afterwards. Confusing the two produces errors with a predictable sign: a monthly response read as a permanent premium overstates the long-run effect, while a long-horizon design read as a monthly response mixes a slow macro trend into an event. This paper separates the two by estimating the response of the green–brown return gap at horizons of zero to twelve months on a single sample and a single control set.

The recent evidence already leans this way. Green stocks outperform brown when climate concern rises unexpectedly, and that outperformance reflects a repricing of existing cash flows rather than a revision of expected returns (Pástor et al., 2022; Ardia et al., 2023). What that evidence does not measure is how long the repricing lasts, which is the question a term structure answers.

The mechanism is not specific to any one market. What it requires is a setting with three features: policy that arrives as dated announcements, so that the timing of the news is observable; green and brown producers inside one market, so that the gap is not contaminated by differences in accounting standards or in the composition of a cross-country index; and a policy regime that shifts within the sample, so that the variation in the index is not the slow drift of a stationary series. Chinese equities satisfy all three, and a horizon structure is a property of the assets rather than of the country.

Climate policy in China has moved in discrete announcements since the dual-carbon target was set in September 2020, and the national emissions trading scheme opened in July 2021. The climate policy uncertainty index of Ma et al. (2023) is constructed by machine-reading six major newspapers with a MacBERT classifier (Cui et al., 2021), so it records the announcement side of policy rather than its implementation. Of the twenty-four months above its ninetieth percentile, sixteen fall in 2018–2022, seven in 2021 alone, and none in 2003–2007. The high-uncertainty months are concentrated where climate policy became salient.

We pair the index with two equity portfolios of firms listed in that market, one green and one brown, and study the gap between their monthly returns. The sample runs from 2015m8 to 2022m12, which yields 88 usable monthly observations after differencing and 76 at the twelve-month horizon. We estimate local projections (Jordà, 2005) of the gap on the standardised monthly change in the index, controlling for the change in the economic policy uncertainty index, at horizons up to twelve months. Inference uses Newey–West standard errors with a bandwidth that grows with the horizon, and a wild bootstrap applied to the design itself supplies horizon-specific critical values.

The results establish three facts. First, a one-standard-deviation increase in climate policy uncertainty raises the gap by 2.67 percentage points in the month of the innovation (t = 2.61, clearing its design-earned threshold of 2.21) and adds nothing thereafter; the twelve-month addition is bounded above by + 0.1 points at ninety per cent. Second, the level of the index is a distinct object whose twelve-month coefficient reaches − 16.9 points on the gap and − 19.2 on the green leg, but no level coefficient clears the threshold its own sampling distribution implies, so the level channel is a description of a pattern rather than an estimate. Third, the five-per-cent threshold is 2.21 for the innovation and 3.97 for the level,

a gap produced by the design rather than the data, and reporting a Newey–West t-statistic against 1.96 at a twelve-month horizon is a statement about the software rather than the coefficient.

This paper pursues four objectives:

- To separate the slow-moving level of a climate-policy uncertainty index from its transitory innovation and to estimate each component's effect on the green–brown gap on one sample, one control set and one set of horizons.
- To replace the conventional 1.96 critical value with thresholds the design earns, by computing the finite-sample distribution of the test itself under a null that holds the regressor path fixed.
- To show that the persistent sign disagreement in the climate-finance literature is a disagreement about the regressor rather than about markets or periods, by reproducing it on a single sample as a competition between the level and the change of the same index.
- To state directly what each coefficient can bear—which are estimates and which are descriptions—and to identify the sample size that would turn the description into an estimate.

Five contributions follow:

- Empirical: we estimate the response of the gap horizon by horizon and show that the addition is concentrated in the announcement month rather than accumulated, which is a different object from the premium studied in long-horizon designs.
- A decomposition and a diagnostic: separating the level from the innovation gives two horizon profiles of opposite slope on the same asset, which reproduces the sign disagreement as a disagreement about the regressor; the diagnostic is available before any structural interpretation.
- A robustness battery that tests the separation against collinearity, orthogonalisation, an error-correction representation, a Beveridge–Nelson decomposition, and the cross-section of asset duration, and reports where each test is silent.
- An inference result: we compute the finite-sample distribution of the test, horizon by horizon, and report the implied thresholds and the size of the test over a grid of persistence and sample length.
- A statement of limits with numbers attached: with twelve-month overlapping windows on 76 observations, the impact estimate is marginal as a family of thirteen, and no level coefficient at an annual horizon clears its own threshold.

The remainder is organised as follows. Section 2 reviews the three literatures the paper draws on and states the gap. Section 3 describes the data and the measurement layer that produces a real-time version of the uncertainty index. Section 4 sets out the methodology: local projections, the level–innovation decomposition, the design-based bootstrap, the finite-sample size correction, and the diagnostic battery. Section 5 reports the results and their discussion. Section 6 concludes and sets out the extensions that would change the answer.

2 Literature Review

Three strands meet in this paper. The first measures uncertainty from text. The second prices green assets. The third estimates dynamic responses at multiple horizons.

2.1 Text-based uncertainty and its effects

The modern measurement of uncertainty from newspapers begins with Baker et al. (2016), who count articles containing terms about the economy, policy and uncertainty, and Bloom (2009), who shows that uncertainty shocks depress investment and hiring before a rebound. Caldara and Iacoviello (2022) apply the same logic to geopolitics. Gentzkow et al. (2019) survey the econometric issues that arise when text becomes a regressor, of which the two that matter here are vocabulary stability and the mapping from article counts to an index scale.

Climate policy uncertainty indices follow this template. Gavriilidis (2021) constructs a US index from eight newspapers; Ma et al. (2023) construct the Chinese index used here from six newspapers with a MacBERT classifier, which handles Chinese word segmentation better than dictionary counts. The Chinese index has two properties that matter for our design. It is monthly, so it can be aligned with return data without interpolation. And it is built from newspapers rather than from official documents, so it records the salience of policy debate rather than the enactment of policy, which is the object that should move prices before implementation.

Two papers have already used this index on Chinese asset prices. Iqbal et al. (2024) estimate a nonlinear ARDL model on daily data and find that a rise in the index lowers carbon allowance prices and raises an ESG index over the long run, while leaving clean energy stocks unaffected. Chen et al. (2024) run an event study over five climate policy events and report ten-day cumulative abnormal returns of −15.4 percent for coal and +11.3 percent for wind power, with the response fading within weeks. The second result bears directly on our question: a market that reprices within weeks is one in which a monthly innovation should show up at impact and should not accumulate. Neither paper separates the level of the index from its innovation.

That distinction is old. In the structural VAR tradition it is the line between permanent and transitory disturbances (Blanchard and Quah, 1989). It returns whenever a text index is used as a regressor, where the level and the first difference are rarely entered side by side, and where the choice between them is made before estimation rather than tested. Section 4.2 is an empirical version of it: we let both enter at the same time and find that they carry opposite signs at annual horizons.

2.2 Green asset pricing

The green asset pricing literature has established three facts. Green bonds trade at a small premium, with estimates clustered between one and nine basis points depending on currency and rating (Zerbib, 2019). Institutional investors with pro-environmental preferences accept lower expected returns in

equilibrium, which implies a higher cost of capital for brown firms (Pástor et al., 2021). And firms with higher carbon emissions earn higher returns, the carbon premium (Bolton and Kacperczyk, 2021).

These results are estimated on annual or multi-year panels. Their identifying variation is cross-sectional, and their object is a long-run difference in expected returns. A monthly innovation in policy uncertainty is not the variation that identifies them, and their estimates do not speak to whether such an innovation moves relative prices. The two designs answer different questions, and the gap between them is where this paper sits.

The closest designs to ours start from that gap. Pástor et al. (2022) show that the recent outperformance of green assets reflects unexpected increases in environmental concern rather than high expected returns: once realised returns are purged of climate-concern shocks, green stocks carry lower expected returns than brown. Ardia et al. (2023) reach the same place from the event side, using a daily media index to show that green stocks rise and brown stocks fall when climate concern increases unexpectedly, and tracing the difference to a fall in the discount rate applied to green firms. For China, Zhao et al. (2024) document a positive greenium in A-shares that widens when public climate concern rises.

The Chinese bond market locates that premium more precisely. Li et al. (2022) compare labelled green bonds with unlabelled bonds that satisfy the same green criteria and find that the discount at issuance exists only for the labelled ones—about twelve basis points, against eight for green bonds pooled—and widens when credit risk is high. A premium that requires a government label prices implicit policy support rather than greenness. If the green advantage in China is conferred by policy rather than by investor taste, then it is a claim on future policy, and the variance of that policy becomes a priced object. Whether that variance helps or hurts green valuations depends on the horizon, which is what the term structure in Section 5 measures.

2.3 Local projections and the term structure of responses

Local projections estimate the response of a variable at horizon h by regressing its h-step-ahead value on the shock and controls (Jordà, 2005). They nest the impulse response of a VAR without imposing its dynamics, and they accommodate nonlinear and state-dependent specifications (Jordà et al., 2013). Plagborg-Møller and Wolf (2021) show that, under correct specification, local projections and VARs estimate the same object.

The cost of the flexibility is inference at long horizons. Overlapping windows induce serial correlation of order h in the residual, which is why we set the Newey–West bandwidth to at least h + 3 and report the effective number of independent observations alongside every long-horizon estimate. With 76 observations at h = 12, that number is about six. The standard remedy is a different asymptotic theory: fixed-b inference treats the ratio of bandwidth to sample size as the object that governs the null distribution (Kiefer and Vogelsang, 2005), and the recommendations of Lazarus et al. (2018) follow from it. We take that route in Section 4.4, where the bootstrap computes the same object by simulation.

Montiel Olea and Plagborg-Møller (2021) show that lag-augmented local projections with normal critical values are uniformly valid across horizons; we report both Newey–West and lag-augmented intervals and show they diverge at the long end in exactly the way the theory predicts.

Table 1
Key studies on uncertainty measurement, green asset pricing and local projections

Study	Method / data	Finding	Gap this paper fills
Baker et al. (2016)	News counts; EPU index	Policy uncertainty is measurable from text	No separation of level from innovation
Ma et al. (2023)	MacBERT news classifier; China CCPU	Monthly climate-policy uncertainty index	Level and change never entered together
Zerbib (2019); Bolton & Kacperczyk (2021)	Long-run panels	Green-bond / carbon premium	No monthly impact response
Pástor et al. (2022); Ardia et al. (2023)	Concern shocks / daily media	Green outperforms brown at impact	No term structure; no level channel
Chen et al. (2024); Iqbal et al. (2024)	Event study / NARDL	Repricing fades within weeks	No level–innovation decomposition
Jordà (2005); Montiel Olea & Plagborg-Møller (2021)	Local projections	Horizon-by-horizon responses	No design-based thresholds for overlapping windows

Notes: The table summarises the literatures reviewed in Sections 2.1–2.3 and locates the contribution of this paper. The split that matters for our result is between studies that treat the climate variable as a persistent level and those that treat it as an innovation or event.

Table 1 shows the evolution from long-run, cross-sectional premium studies to event and daily-media studies, and that none of them enters the level and the innovation of the same index as competing regressors on one sample. This paper fills that gap directly: it shows that the level–innovation split reproduces the literature's sign disagreement on a single sample, which means the disagreement is about the regressor and not about markets or periods.

3 Data and Measurement

This section describes the sources of data used in this study, the layer that produces a real-time version of the uncertainty index, the descriptive statistics of the sample, and the construction of the paired sample. The aim is a well-structured dataset on which the local-projection estimates of Section 5 can be replicated.

3.1 Data sources

The pricing estimates in Section 5 use the final revised CCPU index of Ma et al. (2023). The index is constructed by machine-reading six major Chinese newspapers with a MacBERT classifier and is

released monthly from 2003m4. It records the salience of climate-policy debate rather than the enactment of policy, which is the object that should move prices before implementation.

The green portfolio holds A-share firms whose revenue is classified as environmental protection, renewable energy, energy efficiency or clean technology. The brown portfolio holds firms in coal, thermal power, steel, cement, chemicals and other emission-intensive industries. Both are equal-weighted within month and rebalanced monthly. The return gap, written green-minus-brown (GMB) throughout, is $\mathrm{GMB} = r_{\mathrm{green}} - r_{\mathrm{brown}}$. The classification uses the firms' main business description rather than a third-party rating, so the portfolios can be reconstructed from public filings; the industry mapping is reported in Appendix D.

The macro controls come from the same monthly panel used in Section 3.2. The economic policy uncertainty index is the EPU series of Baker et al. (2016) for China. The market return is the CSI 300. The portfolios are equal-weighted because the return series available to us are aggregated at the portfolio level, with no market capitalisation attached, so a value-weighted version cannot be constructed without firm-level data.

3.2 The real-time nowcast of the CCPU index

Policymakers and investors observe the index with a lag: newspapers are published daily, the index is released monthly, and the underlying articles are revised as databases update. The object relevant to a decision taken in month *t* is the nowcast of the index for month *t* formed with information available in month *t*, not the revised value a researcher observes ex post, because vintages differ and a coefficient estimated on final data need not equal the one a decision-maker could have computed at the time (Giannone et al., 2008; Croushore and Stark, 2001).

We build that nowcast with a dynamic factor model in which the CCPU index enters as an external text block alongside eight monthly indicators for China: CPI inflation, PPI inflation, industrial production, fixed-asset investment, retail sales, the policy interest rate, M2 and an economic-surprise index. The panel covers 2003m4–2022m12, 237 monthly observations after transformation. The index equation excludes the contemporaneous factor, so the index may load on the common state but the common state may not respond to the index within the month; without that restriction the text block would be absorbed into the factor. Estimation is by maximum likelihood through the EM algorithm with the Kalman filter and smoother, which handles the ragged edge of official releases as missing data (Bańbura et al., 2010; Bańbura and Modugno, 2014).

Table 2 reports root mean squared error at four horizons against two benchmarks, an AR(1) and a random walk.

Table 2
Pseudo-real-time nowcast accuracy for the CCPU index (RMSE, index points)

Horizon	Factor model	AR(1)	Random walk	Gain vs AR(1)	Gain vs RW
h = 1	0.464	0.480	0.511	3.4%	9.2%
h = 3	0.567	0.652	0.719	13.0%	21.1%
h = 6	0.574	0.683	0.737	15.8%	22.1%
h = 12	0.594	0.752	0.921	20.9%	35.5%

Notes: Pseudo-real-time evaluation; the model is re-estimated at each origin with the ragged edge preserved.

The short-horizon gain is small, because one month ahead the index is close to a random walk; the gain widens to 21 per cent against the AR(1) and 36 per cent against the random walk at twelve months. At h = 1 an inverse-MSE combination of the factor model and the AR(1) gives 0.454, 11 per cent below the random walk. Diebold and Mariano (1995) tests place the significant gains on the real side of the economy: at h = 12 the factor model beats the AR(1) for industrial production (t = − 2.07), fixed-asset investment (t = − 2.23) and M2 (t = − 2.08). Removing the text block degrades the nowcast of industrial production by up to 11 per cent and leaves prices essentially unchanged.

The measurement layer makes two contributions, neither of which is a pricing estimate. It supplies a real-time object: the regression in Section 5 could be run on any vintage, and the nowcast is the version a market participant would have held at the moment of the announcement. And it validates the index as an economic object: a text index that could not forecast anything would be measuring the newspapers' choice of words, whereas this one predicts the real side of the economy and not prices. It does not enter the pricing regression as a generated regressor, because the revised series is what a replicator can reconstruct and because the difference between vintages enters as classical measurement error that biases the coefficient towards zero—a reliability of 0.9 would put the impact response at 2.97 points and one of 0.8 at 3.34. The estimate is therefore a lower bound in magnitude, conditional on that assumption.

3.3 Descriptive statistics

Table 3 reports the sample. The estimation sample is 2015m8–2022m12, 89 monthly observations, of which 88 enter the regression after differencing. The start date is set by the green portfolio return series, which begin in 2015m8; the brown leg is available earlier and the index itself runs from 2003m4.

Table 3
Descriptive statistics, 2015m8–2022m12

Variable	N	Mean	SD	Min	Max
CCPU index (level)	89	2.445	0.489	1.117	3.627
EPU index	89	269.52	121.54	60.21	661.83
Green portfolio return (%)	89	0.59	9.16	−36.24	22.42
Brown portfolio return (%)	89	−0.18	6.52	−18.13	19.24
Green–brown gap (%)	89	0.76	9.32	−19.96	26.58
CSI 300 return (%)	89	0.02	5.71	−23.62	13.64

Three features of the table drive the results. The gap is more volatile than either portfolio, with a standard deviation of 9.32 percent a month, so a 2.67-point response is 0.29 of a standard deviation, not an outlier-sized move. The green portfolio has a higher mean and a higher variance than the brown portfolio over this window, which is the sample-specific version of the long-run premium debate. And the CCPU index has a standard deviation of 0.49 index points on a mean of 2.45, so a one-standard-deviation innovation is a 20 percent move in the index—a routine month rather than a crisis.

3.4 Sample construction and the end-2022 boundary

The return series are available through 2026m10, 135 months. The CCPU index stops in 2022m12. Because the regression requires both, the paired sample is 89 months. The 46 unused months are a binding constraint on precision, not a choice: extending the index would raise the sample by 52 percent and, at the same signal-to-noise ratio, would raise the t-statistic on the contemporaneous coefficient from about 2.6 to about 3.2. This is the highest-value extension available, and Section 6 states what it would change. We do not attempt it here, because splicing an index built by a different classifier onto the Ma et al. (2023) series would introduce a break exactly in the regressor of interest.

4 Methodology

This section sets out the estimation framework. Section 4.1 defines the local-projection specification. Section 4.2 introduces the level–innovation decomposition. Section 4.3 derives the design-based wild bootstrap and its thresholds. Section 4.4 corrects the finite-sample size of the test. Section 4.5 specifies the diagnostic battery that separates the two components from a mechanical collinearity reading.

4.1 Local projections

The estimating equation is

$$\mathrm{GMB}_{t\to\, t+h} = \alpha_{\,h} + \beta_{\,h} \Delta\, c_t + \gamma_{\,h} \Delta\, u_t + \epsilon_{\,t,h}$$

1

where $\mathrm{GMB}_{t\to\,t+h}$ is the cumulative gap from t to $t+h$, Δc_t is the standardised first difference of the CCPU index, and Δu_t is the standardised first difference of the EPU index. Both regressors are standardised on the estimation sample. The coefficient β_0 is the impact coefficient; β_h, $h \geq 1$, is the cumulative coefficient over h months. Standard errors are Newey–West with bandwidth h + 3 (Newey and West, 1987). Eq. (1) is estimated separately at each horizon, so the path in Table 4 is a sequence of one-shot regressions rather than a single dynamic fit; this is what allows the coefficient to change sign across horizons without the model imposing a shape on it. The design follows the event-study logic of Gürkaynak et al. (2005), who show that asset prices respond to the announcement component of policy rather than only to the action.

4.2 Decomposing level and innovation

The level c_t of the CCPU index and its first difference Δc_t are nearly collinear by construction, and the choice between them is the choice between a persistent and a transitory reading of the same object. We enter both at the same time, on the same sample, so that the two coefficients compete as regressors:

$$\mathrm{GMB}_{t\to\,t+h} = \alpha_h + \beta_h^{L} c_t + \beta_h^{\Delta}\,\Delta c_t + \gamma_h \Delta u_t + \epsilon_{t,h}$$

2

The level coefficient β_h^{L} is the loading of the persistent component; the change coefficient β_h^{Δ} is the loading of the transitory component. The sum $\beta_h^{L} + \beta_h^{\Delta}$ is the response to a permanent one-standard-deviation rise in the index and is the linear combination least exposed to the correlation between the two regressors. Entering both at once converts the literature's sign disagreement into a within-sample competition rather than a between-study one.

4.3 Design-based inference: the wild bootstrap

The threshold that applies to the local-projection coefficient is not 1.96, and computing it is the first step in replacing the conventional critical value with one the design earns. We fit the monthly gap on the EPU change, take the residual, and generate

$$y_t^{*} = \widehat{a} + \widehat{\phi}\, y_{t-1}^{*} + \hat{b}\Delta u_t + e_t v_t$$

3

with v_t a Rademacher variable and $\widehat{\phi} = 0.11$, preserving the heteroskedasticity of the return series. The autoregressive term is there because the monthly residual has first-order autocorrelation of 0.11 and a Ljung–Box(12) statistic of 22.1. Each replication rebuilds the cumulative path and re-estimates equations (1) and (2) at every horizon, so the overlap in the dependent variable is generated rather than assumed away. The regressor path is held at its observed values, which asks how variable the return series alone could make the coefficient—the right question for a design that treats the index as given.

Replications: 2,000. The ninety-fifth percentile of the simulated t-statistic is the design-earned threshold; the bootstrap p-value is the fraction of replications in which the statistic exceeds the observed one.

4.4 Finite-sample size correction

The wild bootstrap of Section 4.3 computes the threshold at the observed sample size, but a second question is how often the nominal five-per-cent test rejects a true null. We answer it by simulation over a grid of regressor persistence and sample length, and by a double bootstrap that takes each replication's own simulated critical value. The fixed-b experiment of Kiefer and Vogelsang (2005) separates the effect of the bandwidth-to-sample ratio from the effect of the sample size: holding b fixed and raising T does not repair the test, because the null distribution depends on b and not on T. This is why a longer sample with the same bandwidth would legitimately report a smaller threshold, and why reporting the ratio rather than the bandwidth is the recommendation of Lazarus et al. (2018).

4.5 The diagnostic battery

The objection the decomposition invites is that a level and its own first difference are mechanically related, so a pair of large opposite coefficients is what collinearity produces. We test this in four ways. The first is the collinearity itself: the correlation, variance inflation factor and condition number of the two standardised regressors at h = 12. The second is orthogonalisation: regressing the change on the level and entering the residual alongside the level, and vice versa. The third is an error-correction representation estimated at the monthly frequency,

$$\Delta \,\mathrm{GMB}_{t+1} = \alpha \; + \phi \,\mathrm{GMB}_{t} + \lambda \, c_t + \beta \, \Delta \, c_t + \gamma \, \Delta \, u_t + \epsilon_{\,t+1}$$

4

which asks whether the gap reverts to a level implied by the index and whether the level channel appears at the one-month horizon. The fourth is a Beveridge–Nelson decomposition (Beveridge and Nelson, 1981), which constructs the permanent and transitory components before any regression is run and therefore does not inherit the level-and-difference parameterisation. A fifth test, the cross-section of asset duration, checks whether the level effect is larger on the long-duration green leg, which is what a discount-rate channel requires and a pure volatility story cannot produce.

5 Results and Discussion

This section reports the estimates, their robustness, the central decomposition, the term-structure interpretation, and the inference limits. Section 5.1 establishes the contemporaneous repricing. Section 5.2 shows it is stable. Section 5.3 is the level–innovation decomposition. Section 5.4 gives the term structure and the two-sector interpretation. Section 5.5 computes what each coefficient can bear.

5.1 The contemporaneous repricing

Table 4 reports the path. The impact coefficient is 2.67 points with a t-statistic of 2.61. Two horizons clear five per cent on the conventional table, h = 0 and h = 1 (t = 1.97), and the second is a cumulative coefficient that contains the first; the bootstrap threshold in the last column, which is the one that applies, is cleared only at h = 0. The path stays positive through eleven months, peaks at 4.94 points at h = 9 and returns to − 0.24 at h = 12.

Table 4
Cumulative and monthly response of the green–brown gap to a one-standard-deviation innovation in the CCPU index

h	β (pp)	NW t	90% interval	Monthly β	Monthly t	p (boot)	N
0	2.67	2.61	[0.98, 4.35]	2.67	2.61	0.020	88
1	3.11	1.97	[0.51, 5.71]	0.44	0.38	0.054	87
2	3.24	1.67	[0.04, 6.44]	0.14	0.14	0.100	86
3	3.03	1.21	[− 1.08, 7.15]	−0.20	−0.20	0.241	85
4	3.97	1.56	[− 0.22, 8.15]	0.86	0.88	0.106	84
5	3.21	1.77	[0.22, 6.19]	−0.84	−0.79	0.053	83
6	1.48	0.90	[− 1.22, 4.18]	−1.61	−1.33	0.348	82
7	2.86	1.42	[− 0.46, 6.19]	1.25	1.04	0.211	81
8	3.87	1.25	[− 1.23, 8.97]	1.01	0.63	0.270	80
9	4.94	1.70	[0.16, 9.71]	1.17	1.63	0.111	79
10	3.55	1.15	[− 1.52, 8.63]	−1.58	−2.08	0.360	78
11	1.80	0.61	[− 3.08, 6.69]	−1.77	−1.83	0.634	77
12	−0.24	−0.10	[− 3.96, 3.49]	0.25	0.27	0.946	76

Notes: Newey–West standard errors, bandwidth h + 3. The cumulative β is the coefficient on the standardised change of the CCPU index in Eq. (1), estimated horizon by horizon. The Monthly β is the single-period response in month t + h, i.e. the same regression with the one-month return as the dependent variable; for h = 0 the two coincide. Intervals are 90 percent and are not adjusted for the thirteen horizons. The last column is the p-value of a wild bootstrap that holds the regressor at its observed path and resamples the monthly return series (4,000 replications); the family-wise values for h = 0 are 0.062 (Romano–Wolf) and 0.119 (Holm).

A persistent premium would produce a coefficient that rises roughly linearly with the horizon, because a monthly difference in expected returns compounds. The estimated path does not: it is 2.67 at impact, 3.11 one month later and − 0.24 at twelve, and the twelve-month interval [− 3.96, 3.49] contains the impact estimate. The sharper test is on the addition rather than on the level. Regressing the sum of the monthly gaps from month one to month twelve on the same regressors gives − 2.61 points (t = − 1.58), with a ninety per cent interval from − 5.32 to + 0.11. This direct-sum regression is close to but not identical to $\beta_{12} - \beta_0$: the two cumulative coefficients differ by − 0.24 − 2.67 = − 2.91, and the 0.3-point gap

is a property of how cumulative local projections define the dependent variable, not a discrepancy. The data cannot reject that the addition is zero, and the upper bound rules out accumulation at the impact rate: eleven further months at 2.67 points would add about 29 points, and the bound is 0.11. The month-by-month increments say the same thing differently: eleven of the twelve are within two standard errors of zero, the exception is month ten (t = − 2.08), and the run of three positive months at seven to nine is the closest the path comes to a trend. The response is non-accumulating rather than zero after impact.

Panel (a) of Fig. 2 makes the point directly. The response in month t + 1 is 0.44 points with a t-statistic of 0.38; in month t + 2 it is 0.14; and the remaining months oscillate around zero with no systematic sign. The entire cumulative path in panel (b) is therefore a step, not a ramp, which is what "front-loaded" means and what "persistent" excludes.

The wild bootstrap establishes the threshold that applies. At impact the nominal five per cent test rejects 8.1 per cent of the time and the implied threshold is 2.21, so the coefficient of 2.61 is a rejection (p = 0.020); a Romano–Wolf stepdown across the thirteen horizons moves the p-value to 0.062, which is marginal rather than comfortable. The impact coefficient is the only one on the path whose single-test p-value is below 0.05, and the monthly responses after the first are indistinguishable from zero. The same bootstrap puts the threshold at 3.97 for the level coefficient at h = 12, against a t-statistic of − 2.31, which is why the level coefficients of Section 5.3 are reported as descriptions rather than rejections.

A 2.67-point move in a series with a monthly standard deviation of 9.32 points is a visible but not extreme month; the sample contains many larger moves in both directions. What is unusual about it is the timing: it is concentrated in the month in which uncertainty rises and does not continue.

5.2 Robustness of the impact estimate

Table 5 reports the impact coefficient under eight alternatives. The estimate is stable in magnitude across all of them; what changes is precision. A market-return control lowers it to 2.19 (t = 2.34), so about four fifths of the raw response is not market beta. A lagged dependent variable leaves it at 2.67 (t = 2.55), which rules out mechanical momentum. Excluding the whole of 2020 raises it slightly to 2.72 (t = 2.64), so the estimate is not a pandemic artefact.

Table 5
Robustness of the contemporaneous (h = 0) coefficient

Specification	β (pp)	NW t	N
Baseline	2.67	2.61	88
+ market return	2.19	2.34	88
+ lagged gap	2.67	2.55	87
+ ΔPPI	2.53	2.61	88
Excluding 2020	2.72	2.64	76
Excluding 2016m1–2017m12	2.28	1.84	64
2018m1–2022m12	2.20	1.70	60
2019m1–2022m12	2.03	1.15	47
2015m8–2018m12	3.32	3.03	40

Notes: Each row is a separate regression of the monthly gap on the standardised change in the CCPU index and the standardised change in the EPU index, with the listed modification. ΔPPI is the standardised change in producer price inflation.

The split sample is the informative part of the table. The first three and a half years give 3.32 points with a t-statistic of 3.03; the last four give 2.03 with a t-statistic of 1.15. The point estimates are close, so the difference is precision rather than sign. But 2015m12 to 2017 marks the supply-side reform, when capacity was cut in coal, steel and cement and the brown portfolio received a large positive shock that had nothing to do with climate policy. Removing those twenty-four months lowers the coefficient to 2.28 (t = 1.84), and restricting the sample to 2018 onward gives 2.20 (t = 1.70). About fifteen per cent of the baseline estimate therefore comes from the reform window, and outside it the estimate is no longer distinguishable from zero at the five per cent level on a single-horizon test. The point estimate is stable—it lies between 2.0 and 3.3 in every row, including the ones that remove the reform—so the finding is not an artefact of one episode. What the reform window supplies is variance, not sign: it is a period in which the index moved and relative returns moved with it, and removing it costs precision.

Figure 3 traces the impact coefficient as the sample is extended and as it rolls. The expanding-window estimate is 3.17 points on the first forty months, falls to a minimum of 1.78 during 2020 and returns to 2.67 on the full sample; its ninety per cent interval contains the full-sample value at every end point. The rolling thirty-six-month estimate is less stable: it peaks at 3.71 in 2018m8, falls to − 0.36 in 2021m7 and recovers to 2.16 by the end of the sample. The picture is an estimate that is stable in level and unstable in any single window.

5.3 Level and innovation are two objects

Table 6 makes the central point directly. At h = 12 we enter the level and the change together, on the same sample.

Table 6
Level and change of the CCPU index at h = 12

Specification	Level (pp)	t	ΔCCPU (pp)	t	N
Level only	−11.86	−1.97	—	—	76
Change only	—	—	−0.94	−0.44	76
Both	−16.86	−2.30	8.10	1.90	76
Both + ΔEPU	−16.88	−2.31	8.90	2.06	76
Both + market	−16.72	−2.30	7.36	1.66	76
Both + trend	−21.23	−3.61	10.52	2.69	76

Notes: All rows use the same 76 observations. "Change only" omits the ΔEPU control that enters Eq. (1); with that control the coefficient is − 0.24, which is the h = 12 entry of Table 4. "Both" is the CCPU level and CCPU change; "Both + ΔEPU" is the S7 specification used in Sections 5.4, 5.5 and the placebo of Section 5.5. "Level only" omits the EPU change, so its − 11.86 differs from the CCPU level alone of Table 9, which includes it. "Market" is the cumulative CSI 300 return over the same window; "trend" is a standardised linear trend.

The two coefficients carry opposite signs and the level coefficient is the larger in absolute value. Adding a market return leaves both essentially unchanged, so the pattern is not a beta effect. Adding a linear trend strengthens both, which is what happens when a regressor is correlated with the trend and the trend is not the mechanism. Adding the EPU change moves the change coefficient from 8.10 to 8.90 and leaves the level coefficient where it was. This is the paper's central decomposition: the CCPU index, used at the monthly frequency, contains a persistent component and a transitory component, and the two are associated with the green–brown gap with opposite signs. The transitory component is statistically distinguishable from zero and the level component is a description rather than an estimate. A researcher who wants the effect of an announcement should use the transitory component, which is what the impact estimate in Table 4 does. A researcher who wants the effect of living in an uncertain policy environment should use the persistent component. The two questions have different answers, and the literature's disagreement follows from asking them with the same regressor.

Table 7 makes that mapping explicit. It sorts the prior literature on climate and green–brown pricing by the form in which the climate variable enters, and shows that the sign disagreement follows the form rather than the market.

Table 7
How the climate variable enters prior studies, and the sign it produces

Study	Climate variable form	Green–brown finding	Reading
Zerbib (2019)	Level (bond yield, long run)	Green-bond premium + 1 to + 9 bp	Level, premium reading
Bolton & Kacperczyk (2021)	Level (cross-section of expected return)	Carbon premium: brown earns more	Level, premium reading
Pástor et al. (2021)	Level (equilibrium expected return)	Lower expected return for green	Level, premium reading
Li et al. (2022)	Level (at-issuance price)	Green-bond discount − 12 bp	Level, discount reading
Zhao et al. (2024)	Concern level / change	Greenium widens as concern rises	Level, premium reading
Engle et al. (2020); Barnett et al. (2020)	Level (persistent uncertainty state)	Uncertainty priced as a persistent state	Level, persistent-state reading
Pástor et al. (2022)	Innovation (purified concern shocks)	Green outperforms at impact; residual expected return lower	Innovation
Ardia et al. (2023)	Innovation (unexpected daily concern)	Green up, brown down at impact	Innovation
Chen et al. (2024)	Innovation (event, 10-day)	Wind + 11.3%, coal − 15.4%	Innovation
Iqbal et al. (2024)	Level (long-run NARDL, daily)	ESG up, carbon allowance down	Level, mixed

Notes: The split is between studies that treat the climate variable as a persistent level and those that treat it as an innovation or event. The level studies disagree among themselves on sign because the level mixes a premium reading with a discount reading; the innovation studies agree that green outperforms brown at impact. The paper reproduces this disagreement on a single sample by entering the level and the change of the same index as competing regressors (Table 6).

The objection this decomposition invites is that the two regressors are a level and its own first difference and are therefore mechanically related. The objection is testable, and the four diagnostics of Section 4.5 rule it out rather than assume it away.

The first is the collinearity itself. On the 76 observations used at h = 12, the correlation between the two standardised regressors is 0.58, the variance inflation factor is 1.50 for each, and the condition number of their correlation matrix is 3.71. A variance inflation factor of 1.50 inflates each standard error by 22 per cent, not by a factor of five. The joint Wald test that both coefficients are zero gives 5.43 on two degrees of freedom (p = 0.066). And the restriction that the two coefficients are equal and opposite, which is the mechanical reading, is rejected: the sum is − 8.8 points with a t-statistic of − 2.19 (p = 0.032).

The second is orthogonalisation. Regressing the change on the level and entering the residual alongside the level leaves the level coefficient at − 11.9 (t = − 2.32); the point estimate falls by five points relative to the joint regression, which is what the shared variation accounts for, and it stays negative and significant. Regressing the level on the change and entering that residual alongside the change leaves the change coefficient at 8.1 (t = 1.90). Neither coefficient survives only because the other is present.

Figure 5 shows what these first two tests are about. The confidence region in panel (a) is elongated along the negative diagonal, which is the visual signature of two correlated regressors, and it is bounded away from the locus on which the coefficients are equal and opposite; the restriction is rejected at the five per cent level. Panels (b) and (c) show the slopes in the scatter: the level coefficient is the downward slope in panel (b), the change coefficient is the upward slope in panel (c), and neither is carried by a handful of points.

The third is an error-correction representation. The adjustment coefficient ϕ is 0.089 with a t-statistic of 1.22, which is the wrong sign for error correction and insignificant: the gap does not revert to a level implied by the index. The level coefficient λ is + 1.23 (t = 1.13) and the change coefficient is + 3.33 (t = 2.49). Read alone, this looks like a rejection of the level channel; it is not. The level coefficient in the cumulative regression is + 1.40 at h = 0 and + 1.15 at h = 1, and it turns negative only from h = 2. A specification that asks what happens over the next month cannot recover a level effect that has not yet appeared.

The fourth is a Beveridge–Nelson decomposition, which constructs the two components before any regression is run. At h = 12 the permanent component enters at − 15.8 (t = − 2.24) and the transitory component at + 3.0 (t = 0.97). At h = 9 both are significant and of opposite sign: −12.2 (t = − 2.43) and + 7.7 (t = 2.37). The decomposition supports the permanent half of the reading firmly and the transitory half only at the nine-month horizon.

Taken together, the four diagnostics and the Beveridge–Nelson decomposition rule out the mechanical reading. They do not rule out that the level coefficient is the loading of a slow-moving omitted factor for which two persistent indices are proxies; Section 5.5 states the consequence.

5.4 The term structure and the duration channel

The decomposition in Table 6 is Eq. (2) evaluated at every horizon, with the change and the EPU change as controls throughout. It has a shape.

Table 8 reports the two coefficient paths.

Table 8
Level, change and sum coefficients by horizon

h	Level	t	L thr	Change	t	Δ thr	Sum	t	Sum thr
0	1.40	1.22	2.58	1.93	1.84	2.20	3.33	2.57	2.31
1	1.15	0.65	2.79	2.50	1.46	2.34	3.65	1.91	2.42
2	−0.19	−0.08	2.94	3.34	1.41	2.42	3.15	1.41	2.53
3	−1.79	−0.53	3.08	3.98	1.21	2.48	2.19	0.81	2.62
4	−3.73	−0.91	3.13	5.97	1.70	2.57	2.24	0.76	2.71
5	−5.97	−1.40	3.16	6.43	2.24	2.56	0.46	0.18	2.85
6	−5.72	−1.26	3.22	4.54	1.81	2.69	−1.18	−0.41	2.81
7	−4.66	−0.99	3.22	5.44	2.07	2.70	0.78	0.23	2.82
8	−6.56	−1.31	3.20	7.46	2.00	2.71	0.90	0.22	2.87
9	−10.64	−1.91	3.24	10.67	2.59	2.99	0.03	0.01	2.70
10	−14.68	−2.66	3.65	11.59	2.54	3.27	−3.09	−0.84	3.05
11	−17.84	−2.68	3.87	11.58	2.13	3.28	−6.25	−1.96	3.44
12	−16.88	−2.31	3.97	8.90	2.06	3.63	−7.97	−2.04	3.42

Notes: Eq. (2) estimated separately at each horizon. Level and Change are the coefficients on the standardised level and the standardised change of the CCPU index; the control is the standardised change of the EPU index. Newey–West standard errors, bandwidth h + 3. Coefficients in percentage points. L thr, Δ thr and Sum thr are the ninety-fifth percentiles of the wild-bootstrap null distributions (the sum is $\beta_h^L + \beta_h^\Delta$).

The level path has a shape. The coefficient is + 1.40 at impact and indistinguishable from zero; it crosses zero at two months; it reaches − 6.0 at five months, − 10.6 at nine and − 17.8 at eleven. The change path moves the other way: +1.93 at impact, + 6.4 at five and + 11.6 at eleven. The two are not two estimates of one object with opposite noise; they are two profiles with opposite slopes. From the third month onward the five per cent threshold for the level coefficient lies between 3.0 and 4.0, and the largest t-statistic on the path, − 2.68 at eleven months, is two thirds of the threshold that applies to it. No point on this path is a rejection.

The sum is the exception, at one end. Because it is a linear combination along the direction in which the two regressors are well identified, its null distribution is tighter: the threshold is 2.31 at impact and 3.42 at twelve months. The sum therefore rejects at impact (3.33, t = 2.57, p = 0.034) and does not reject at twelve months (− 7.97, t = − 2.04, p = 0.257). A permanent one-standard-deviation rise in the index is

associated with higher relative valuations when it happens and with a lower relative valuation a year later; the first half of that sentence is an estimate and the second half is a description.

The cross-section says the same thing more sharply. Table 9 repeats Eq. (2) on each portfolio leg separately.

Table 9
Level and change coefficients by portfolio leg, at selected horizons

h	Portfolio	Level	t	threshold	Change	t
0	Green	0.66	0.69	—	1.72	1.32
0	Brown	−0.74	−0.74	—	−0.21	−0.26
3	Green	−1.70	−0.50	—	3.55	1.23
3	Brown	0.09	0.05	—	−0.43	−0.30
6	Green	−5.90	−1.30	—	4.42	1.64
6	Brown	−0.18	−0.07	—	−0.13	−0.06
9	Green	−11.08	−1.82	—	9.84	2.93
9	Brown	−0.44	−0.12	—	−0.83	−0.30
12	Green	−19.18	−4.17	4.19	10.91	4.12
12	Brown	−2.31	−0.42	3.81	2.01	0.57

Notes: Each row is a separate regression of the cumulative return of that leg on the same regressors as Table 8. Coefficients in percentage points.

The level effect is a green-side effect. The coefficient is − 19.2 on the green leg (t = − 4.17) and − 2.3 on the brown one (t = − 0.42); the standard error on the green leg is 4.60 against 7.32 on the gap. The ordering is what a duration channel requires—the long-lived asset carries the discount-rate exposure and the short-lived one does not. The green leg's level coefficient sits on the boundary of its own threshold (bootstrap p = 0.051), which is a borderline rejection, not a rejection. The difference between the two legs, estimated directly by the gap regression as $\beta_G - \beta_B = -16.9$ (t = − 2.31), does not clear its own threshold of 3.97 (bootstrap p = 0.22), so the ordering across legs is a description rather than a rejection. It also survives the obvious objection, which is that the green portfolio is simply more volatile: a pure volatility story would give the two legs coefficients of the same sign at every horizon, and the green level coefficient changes sign between impact and twelve months while the brown one does not move at all.

A two-sector model organises the two signs. Green capital is long-lived and costly to reverse; brown capital is short-lived and closer to reversible. Duration makes the green portfolio more exposed to the discount-rate component of an innovation in uncertainty, which is negative and increasing in the horizon. The option to delay, and the option value of being the technology that policy substitutes towards, make it

more exposed to the arrival component, which is positive and decays. Investors discount at $r + \pi$, where π is an uncertainty premium that rises with the perceived variance of future policy, and the price of sector i is

$$P_i(\pi) = \sum_{t=1}^{T_i} \frac{y_i}{(1 + r + \pi)^t}$$

5

so the semi-elasticity of the price with respect to the premium is the negative of the Macaulay duration, $-D_i$, and $D_G > D_B$. A persistent rise in π therefore lowers the green price more than the brown price. The same duration logic organises the cross-section of equity returns: Lettau and Wachter (2007) show that growth firms behave like long-duration assets and are more sensitive to discount-rate news. A one-off rise in uncertainty does something else: the value of the option to delay and to be the replacement technology scales with the square root of the horizon, so the impact repricing is

$$\Delta \ln P_i = -D_i \Delta \pi + \kappa \sqrt{T_i} \Delta \sigma$$

6

with $\kappa > 0$. The first term is negative and increasing in T_i; the second is positive and also increasing in T_i. The sign of the total effect depends on which term dominates at the horizon being measured.

Figure 6 plots the two terms against the horizon, with the estimates superimposed. Panel (a) shows the level channel: a hundred-basis-point rise in the premium lowers the price by an amount that rises with T, from about 2.8 percent at five years to about 8.5 percent at twenty. The curve is scaled so that the twenty-year asset matches the green estimate of − 19.2 points; the five-year point is then a prediction, and the brown estimate sits above it at − 2.3 with an interval that contains the prediction. Panel (b) shows the impact channel, scaled so that the twenty-year asset matches the green impact estimate; the five-year prediction is + 1.03, and the brown estimate is − 0.60 with an interval whose upper bound is 0.61. The model over-predicts the short-duration impact response, and on the impact channel it gets the sign of the cross-sectional difference wrong: it predicts a positive impact for brown and the data give a small negative one. The model explains the half of the decomposition that the paper rests on, and fails on the half it does not.

The level coefficient converts into a required return, and the conversion is a check on whether the duration channel is quantitatively plausible. At a discount rate of five per cent, a twenty-year annuity has a Macaulay duration of 8.48 years, so a hundred-basis-point rise in the premium lowers its price by 8.48 per cent. The green coefficient of − 19.18 points therefore corresponds to a premium rise of 226 basis points; the same calculation for the brown leg converts − 2.31 points into 84 basis points. At 226 basis points this is of a different order of magnitude from the twelve-basis-point discount that Li et al. (2022) find for labelled green bonds and from the premium cluster of one to nine basis points that Zerbib (2019)

documents. The cross-section is less extreme than the raw coefficients suggest, because the durations differ by a factor of three; the green-to-brown ratio is 0.12 in the data against 0.33 in the model. The magnitude is where the two readings—a discount-rate channel and a slow-moving omitted factor—differ, and it is the reason the omitted factor is treated as the first limit in Section 5.5.

A reader coming from Ardia et al. (2023) will notice that our level channel runs against their discount-rate result. They find that the discount rate applied to green firms falls when climate concern rises; ours measures uncertainty about what climate policy will be enacted, a shift in the distribution of future payoffs, which raises the premium demanded for holding capital whose value depends on that policy. A preference shock and a risk shock can move the green discount rate in opposite directions, and our design does not separate them: the index contains both, and the level coefficient is the net of the two. What the estimate establishes is which component dominates at the annual horizon in this sample—the risk component—and the net level effect is negative and large.

5.5 Inference limits and what the coefficients can bear

This section computes the precision the design actually has, because nearly every t-statistic in the paper comes from a regression on an overlapping window with a Newey–West standard error whose bandwidth grows with the horizon. The question is how often that statistic would exceed 1.96 if the index had no effect at all, and what number should replace 1.96.

Table 10 reports the rejection rate at the nominal five per cent level and the threshold implied by the simulated distribution, with the change and with the level as the regressor of interest.

Table 10
Size and implied threshold of the local-projection test, nominal level five per cent

h	Size, change	Threshold, change	Size, level	Threshold, level
0	8.1	2.21	11.2	2.48
2	5.0	1.96	15.4	2.76
4	3.1	1.82	18.5	2.97
6	1.8	1.68	18.6	2.98
8	4.5	1.92	19.5	3.07
10	5.8	2.03	25.3	3.26
12	8.7	2.23	27.3	3.45

Notes: 2,000 wild-bootstrap replications per horizon, Rademacher weights, seed 20261104. Size is the fraction of replications in which the Newey–West t-statistic exceeds 1.96 in absolute value when the null is true; the threshold is the ninety-fifth percentile of that statistic. The regressor path is held at its observed values. Entering the change beside the level at h = 12 raises the threshold to 3.97.

The size of the test is not a constant, and it moves in both directions. At impact the change design rejects 8.1 per cent of the time, so 1.96 is mildly too generous and the threshold is 2.21. Between four and six months it rejects 1.8 to 3.1 per cent, so the same 1.96 is too strict and the test loses power exactly where a persistent premium would show up. At twelve months with the level as the regressor it rejects 27.3 per cent of the time, and the threshold that would deliver five per cent is 3.45. A paper that reports Newey–West t-statistics at every horizon against one critical value is therefore reporting three different tests under one name.

The impact coefficient is tested. Its t-statistic of 2.61 against a threshold of 2.21 gives a bootstrap p-value of 0.020, and the coefficient survives the correction the design requires. Power at the estimated effect is about 62 per cent; to be detected 80 per cent of the time the effect would have to be about 1.2 times larger, roughly 3.3 percentage points. No level coefficient at an annual horizon is tested. At h = 12 the coefficient is − 2.31 against a threshold of 3.45 on its own and 3.97 beside the change; at h = 10 and h = 11 the t-statistics are − 2.66 and − 2.68 against 3.65 and 3.87. Every point on that path falls short of the number that applies to it, and the path taken as a whole falls inside its simultaneous band, whose critical value is 4.79. The weighted slope of the path is − 1.40 points a month, against a null ninety per cent interval of − 1.50 to + 1.48 and a p-value of 0.12. The level channel is a description of a shape, and the shape is suggestive without being established. The green leg's level coefficient is − 19.18 with a t-statistic of − 4.17 and a threshold of 4.19 (p = 0.051): a borderline rejection. The sum of the level and change coefficients rejects at impact (p = 0.034) and does not reject at twelve (p = 0.257). That asymmetry is the paper's central inference result: the same permanent rise is priced when it happens and cannot be shown to be priced a year later.

Three candidates explain the variation in Table 10: the persistence of the regressor, the number of observations, and the bandwidth. Table 11 separates them by simulation, holding the design otherwise fixed at a twelve-month window.

Table 11
Threshold and rejection rate at h = 12 by persistence and sample length

Persistence ρ	T = 76	T = 150	T = 300
0.0	2.62 (15.5; 6.2)	2.27 (8.3; 5.5)	2.07 (8.0; 6.5)
0.4	2.92 (18.5; 4.0)	2.44 (14.5; 7.5)	2.20 (9.2; 5.0)
0.7	3.48 (26.5; 6.2)	2.72 (14.2; 3.0)	2.40 (11.3; 5.5)
0.9	4.58 (31.8; 2.2)	3.22 (19.3; 5.0)	2.69 (13.8; 5.5)

Notes: Each cell is the ninety-fifth percentile of the simulated t-statistic; the two numbers in parentheses are rejection rates against a nominal five per cent, the first with a fixed 1.96 critical value and the second with a double bootstrap. 400 outer and 200 bootstrap replications per cell, seed fixed within cell. The bandwidth is h + 3 throughout, so b = 15/T varies across columns by construction.

Persistence matters, and it is not the whole story. Moving ρ from zero to 0.9 at T = 76 raises the threshold from 2.62 to 4.58. But at ρ = 0, where the regressor is white noise, the test still rejects 15.5 per cent of the time: with a twelve-month window on 76 observations and a bandwidth of 15, the overlap alone is enough to break the normal approximation. The estimated persistence of the index in the asset sample is 0.41. The double bootstrap repairs the rate in every cell, to between 2.2 and 7.5 per cent with an average of 5.2. Holding b = 0.20 and raising T from 76 to 300, the rejection rate goes from 22.0 to 15.5 to 16.5 per cent and the threshold from 2.92 to 2.75 to 2.64: the extra observations do not repair the test, because the null distribution depends on b and not on T. The lag-augmented design is not the conservative alternative it is often taken to be: with twenty-nine parameters on 76 observations at h = 12, its heteroskedasticity-robust standard error is biased downward and it rejects a true null 18.1 per cent of the time at h = 7.

Figure 4 The twelve-month coefficient across ten specifications, grouped by the regressor whose coefficient is reported. Bars are 90 percent Newey–West intervals; filled markers are those with a t-statistic beyond 1.96 in absolute value, hollow markers those below it, a threshold that Section 5.5 shows to be too generous for the level at this horizon. All ten specifications use the same 76 observations at h = 12. The dashed horizontal line separates the seven specifications that report a coefficient on the level of the index from the three that report a coefficient on its change.

The range is wide, and its width is entirely in one group. Across the seven level specifications the coefficient runs from − 17.7 to + 3.0 points, and six of the seven are negative with a t-statistic beyond 1.96 in absolute value; Section 5.5 shows that this threshold is too generous for the level at this horizon, so the six signs are a pattern and not six rejections. Across the three change specifications it runs from − 0.9 to + 1.5 and none is distinguishable from zero. The most negative specification is the level with the raw EPU control and a lagged gap, at − 17.7 points (t = − 3.90). If the ten coefficients were ten draws on one object, that separation would not appear. The disagreement across specifications is the signature of two different regressors, not of an unstable estimate of one, and a paper reporting one of these ten numbers is reporting its choice of control set.

A test that rejects for any persistent regressor would be worthless, so the level coefficient needs a placebo. We ran two. The first permutes the index: we cut the CCPU series into twelve-month blocks and shuffled the blocks 500 times, leaving the return series untouched, and re-estimated the h = 12 level coefficient under the S7 specification, which returns the observed − 16.88. A persistent series with this autocorrelation but no time alignment with the gap produces a distribution centred near zero, and the observed value lies below its fifth percentile (− 11.73); only 0.6 per cent of the permutations produce a value that low. The second replaces the index with the general policy-uncertainty index. Table 12 reports the two indices entered together.

Table B2
Lag-augmented local projections

h	β (pp)	Robust se	t	Parameters
0	3.18	1.17	2.72	5
1	4.31	1.67	2.57	7
2	4.68	2.24	2.09	9
3	4.26	2.83	1.51	11
4	4.43	3.27	1.36	13
5	3.88	3.27	1.19	15
6	2.46	3.71	0.66	17
7	4.33	4.29	1.01	19
8	6.60	4.90	1.35	21
9	4.12	4.82	0.86	23
10	−3.12	4.72	−0.66	25
11	−5.54	4.80	−1.15	27
12	−8.39	4.37	−1.92	29

Notes: Each row adds h + 1 lags of Δc and Δu to Eq. (1). The parameter count includes the intercept.

Table 12
The level coefficient at h = 12 with both indices entered

Regressor	Coefficient (t)	Threshold	p (boot)
CCPU level only	−12.25 (− 2.00)	3.45	0.263
EPU level only	20.53 (4.69)	—	—
CCPU level, with EPU level	−16.43 (− 3.51)	4.00	0.083
EPU level, with CCPU level	23.28 (5.68)	4.79	0.027

Notes: Cumulative gap over twelve months, 76 observations. The first two rows are separate regressions with the EPU change as control; the first row's coefficient (− 12.25) includes the EPU change, which is why it differs from the − 11.86 of Table 6 'Level only' (which omits it). The last two rows are one regression containing both levels and the EPU change; their thresholds come from the joint two-index null.

The two indices carry opposite signs, and the sign of each survives the presence of the other. Climate policy uncertainty is associated with a discount on the green leg that grows with the horizon; general policy uncertainty is associated with the reverse. The economic reading is that they are not two measures of one thing: general policy uncertainty is a brown-side risk, because brown firms are the ones

whose subsidies, permits and input prices are set by policy, while climate policy uncertainty is a green-side risk, because the green advantage in this market is itself conferred by policy (Li et al., 2022). A single "uncertainty" regressor averages the two and can carry either sign depending on which dominates in the sample, which is one more reason the literature disagrees. With both levels in the regression the climate coefficient keeps its sign and gains precision, but only the general index clears its threshold; the climate coefficient reaches p = 0.083. The level channel is therefore identified as a sign and as an ordering across legs, not as a magnitude.

Five limits qualify the estimates, and each has a specific consequence. An omitted persistent factor is the limit that matters: a slow-moving variable that raises measured climate policy uncertainty and lowers long-duration valuations would produce the level coefficient without any discount-rate channel, and 226 basis points of implied premium is large enough that the alternative cannot be dismissed on plausibility. The market control and the Beveridge–Nelson decomposition rule out a factor priced in the equity market or absorbed by the permanent-transitory split of this index; they do not rule out a factor outside both, and the cross-section narrows but does not close the field. The estimates are conditional associations rather than causal responses: a regression of the index on six lags of industrial production rejects the null of no predictability with F = 5.02, so months in which uncertainty rises are not randomly assigned. Precision at the long end produces a division of labour rather than a caveat: the impact coefficient is an estimate, the level coefficients are a description, and only the green leg's − 19.2 sits close enough to its threshold to be read as a borderline rejection. Multiple testing is handled by the design-earned wild-bootstrap threshold of Section 4.3, under which the impact coefficient clears at a bootstrap p-value of 0.020; Romano–Wolf gives 0.062 and Holm gives 0.119, so the rejection does not survive the most conservative family-wise corrections, but the estimate is unaffected and the verdict is qualified, not reversed. The portfolios are equal-weighted because the return series carry no market capitalisation, and the classification is fixed at the 2022 vintage; a value-weighted replication would answer a different question and cannot be run with these data.

6 Conclusion and Future Work

A climate-policy innovation reprices the green–brown gap by 2.67 points in the month it arrives, 0.29 of the gap's monthly standard deviation, and adds nothing afterwards: the total addition over the following twelve months is − 2.61 points, with a ninety per cent interval from − 5.32 to + 0.11. The level of the index is a second object. It has no impact effect and a coefficient that reaches − 16.9 points at twelve months on the gap and − 19.2 on the green leg alone, and none of those coefficients clears the threshold that its own sampling distribution implies. The two objects coexist because duration enters the discount-rate channel and the option-value channel with opposite signs.

The contribution of this paper is the demonstration that one regressor cannot recover both, that the disagreement in the climate-finance literature is reproducible on a single sample as a disagreement about the regressor rather than about markets or periods, and that the sample required to test the second object is larger than the one available. The ten specifications in Fig. 4 make the first point: the

seven that report a coefficient on the level span twenty points while the three that report a coefficient on the change sit within two points of zero. A single long-horizon number is a choice of control set.

The inference result travels with the application. On this design the five per cent threshold is 2.21 for the impact coefficient and 3.97 for the level coefficient at twelve months, and the gap between them is a property of the design rather than the data: a grid over persistence and sample length shows the nominal test rejecting a true null 8 to 32 per cent of the time at twelve months, with the rejection rate governed by the ratio of bandwidth to sample rather than by the persistence of the regressor. Reporting a Newey–West t-statistic against 1.96 at a twelve-month horizon is a statement about the software, not about the coefficient.

For policy the implication is direct. An innovation is associated with a movement in relative valuations in the month it arrives and leaves them where they were a year later, so a policy that surprises the market buys a month of relative green outperformance and no more. What persists is the level: a standard deviation of the index is associated with 226 basis points of implied required return on long-duration green capital, and the conversion is a falsification check rather than a measured effect. Stable climate policy is associated with a lower cost of capital for green investment through the level of uncertainty, not through the frequency of announcements, and a strategy of announcing more and legislating less would show up in the first number and not in the second.

Two extensions would change the answer rather than extend it. Carrying the index to 2026 would raise the paired sample by 52 per cent and the impact t-statistic from about 2.6 to about 3.2, which is enough to test the long end instead of describing it; at the larger sample the twelve-month level t-statistic would move from − 2.31 toward the 3.97 threshold that currently leaves it a description, so the level channel would become an estimate rather than a shape. A provincial panel with shift-share exposure to national policy would supply the exogenous variation this design lacks, and would let the level coefficient be read as a response rather than an association. Within this sample the impact repricing is measurable, its horizon is one month, the level effect has a term structure that the cross-section of durations orders correctly, and the threshold that separates the two is computable before the estimates are read.

Declarations

Competing interests

The author has no competing interests to declare that are relevant to the content of this article.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contribution

R.L. conceived the research question, constructed the data, carried out the estimation and inference, produced all tables and figures, and wrote the manuscript.

Data Availability

The climate policy uncertainty index is from Ma et al. (2023). The macro panel and the equity portfolios are constructed from public sources; the replication package described in Appendix E contains the scripts, the derived series and the seeds needed to reproduce every table and figure.

Code availability

The pricing regressions, the diagnostics and the figures are computed in Python 3.11; the nowcast of Section 3.2 uses the dynamic factor model estimated in MATLAB R2023b (EM algorithm and Kalman smoother). Scripts are listed in Appendix E.

References

1. Antolín-Díaz, J., Drechsel, T., & Petrella, I. (2024). Advances in nowcasting economic activity: The role of heterogeneous dynamics and fat tails. *Journal of Econometrics*, *238*(2). https://doi.org/10.1016/j.jeconom.2023.105634. Article 105634.
2. Ardia, D., Bluteau, K., Boudt, K., & Inghelbrecht, K. (2023). Climate change concerns and the performance of green vs. brown stocks. *Management Science*, *69*(12), 7607–7632. https://doi.org/10.1287/mnsc.2022.4636
3. Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, *131*(4), 1593–1636. https://doi.org/10.1093/qje/qjw024
4. Bańbura, M., Giannone, D., & Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, *25*(1), 71–92. https://doi.org/10.1002/jae.1137
5. Bańbura, M., & Modugno, M. (2014). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of Applied Econometrics*, *29*(1), 133–160. https://doi.org/10.1002/jae.2306
6. Barnett, M., Brock, W., & Hansen, L. P. (2020). Pricing uncertainty induced by climate change. *Review of Financial Studies*, *33*(3), 1024–1066. https://doi.org/10.1093/rfs/hhz144
7. Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. *Operational Research Quarterly*, *20*(4), 451–468. https://doi.org/10.1057/jors.1969.103
8. Beveridge, M. E., & Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components, with particular attention to measurement of the

business cycle. *Journal of Monetary Economics*, *7*(2), 151–174. https://doi.org/10.1016/0304-3932(81)90040-4

9. Blanchard, O. J., & Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review*, *79*(4), 655–673.
10. Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, *77*(3), 623–685. https://doi.org/10.3982/ECTA6248
11. Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, *142*(2), 517–549. https://doi.org/10.1016/j.jfineco.2021.05.008
12. Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, *112*(4), 1194–1225. https://doi.org/10.1257/aer.20191823
13. Campbell, J. Y., & Ammer, J. (1993). What moves the stock and bond markets? A variance decomposition for long-term asset returns. *Journal of Finance*, *48*(1), 3–37. https://doi.org/10.1111/j.1540-6261.1993.tb04700.x
14. Chen, D., Zeng, Z., & Chen, Y. (2024). Heterogeneous impacts of multiple climate policies on the Chinese stock market. *Finance Research Letters*, *60*., Article 104816. https://doi.org/10.1016/j.frl.2023.104816
15. Croushore, D., & Stark, T. (2001). A real-time data set for macroeconomists. *Journal of Econometrics*, *105*(1), 111–130. https://doi.org/10.1016/S0304-4076(01)00072-0
16. Cui, Y., Che, W., Liu, T., Qin, B., Wang, S., & Hu, G. (2021). Revisiting pre-trained models for Chinese natural language processing. In *Findings of the Association for Computational Linguistics: EMNLP 2020*(pp. 6578–6588). https://doi.org/10.18653/v1/2020.findings-emnlp.586
17. Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, *13*(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599
18. Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press. https://doi.org/10.1515/9781400830176
19. Doz, C., Giannone, D., & Reichlin, L. (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics*, *164*(1), 188–205. https://doi.org/10.1016/j.jeconom.2011.02.012
20. Engle, R. F., Giglio, S., Kelly, B., Lee, H., & Stroebel, J. (2020). Hedging climate change news. *Review of Financial Studies*, *33*(3), 1184–1216. https://doi.org/10.1093/rfs/hhz072
21. Gavriilidis, K. (2021). Measuring climate policy uncertainty. SSRN Working Paper No. 3847388. https://doi.org/10.2139/ssrn.3847388
22. Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. *Journal of Economic Literature*, *57*(3), 535–574. https://doi.org/10.1257/jel.20181020
23. Giannone, D., Reichlin, L., & Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, *55*(4), 665–676. https://doi.org/10.1016/j.jmoneco.2008.05.010

24. Gürkaynak, R. S., Sack, B., & Swanson, E. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, *1*(1), 55–93.
25. Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, *6*(2), 65–70.
26. Hong, H., Li, F. W., & Xu, J. (2019). Climate risks and market efficiency. *Journal of Econometrics*, *208*(1), 265–281. https://doi.org/10.1016/j.jeconom.2018.09.015
27. Iqbal, N., Bouri, E., Shahzad, S. J. H., & Alsagr, N. (2024). Asymmetric impacts of Chinese climate policy uncertainty on Chinese asset prices. *Energy Economics*, *133*., Article 107518. https://doi.org/10.1016/j.eneco.2024.107518
28. Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182. https://doi.org/10.1257/0002828053828518
29. Jordà, Ò., Schularick, M., & Taylor, A. M. (2013). When credit bites back. *Journal of Money Credit and Banking*, *45*(s2), 3–28. https://doi.org/10.1111/jmcb.12069
30. Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *Review of Financial Studies*, *33*(3), 1067–1111. https://doi.org/10.1093/rfs/hhz137
31. Kiefer, N. M., & Vogelsang, T. J. (2005). A new asymptotic theory for heteroskedasticity-autocorrelation robust tests. *Econometric Theory*, *21*(6), 1130–1164. https://doi.org/10.1017/S0266466605050565
32. Lazarus, E., Lewis, D. J., Stock, J. H., & Watson, M. W. (2018). HAR inference: Recommendations for practice. *Journal of Business & Economic Statistics*, *36*(4), 541–559. https://doi.org/10.1080/07350015.2018.1506344
33. Lettau, M., & Wachter, J. A. (2007). Why is long-horizon equity less risky? A duration-based explanation of the value premium. *Journal of Finance*, *62*(1), 55–92. https://doi.org/10.1111/j.1540-6261.2007.01201.x
34. Li, Q., Zhang, K., & Wang, L. (2022). Where's the green bond premium? Evidence from China. *Finance Research Letters*, *48*., Article 102950. https://doi.org/10.1016/j.frl.2022.102950
35. Ma, Y. R., Liu, Z., Ma, D., Zhai, P., Guo, K., Zhang, D., & Ji, Q. (2023). A news-based climate policy uncertainty index for China. *Scientific Data*, *10*(1). https://doi.org/10.1038/s41597-023-02817-5. Article 881.
36. Marcellino, M., & Schumacher, C. (2010). Factor MIDAS for nowcasting and forecasting with ragged-edge data: A model comparison for German GDP. *Oxford Bulletin of Economics and Statistics*, *72*(4), 518–550. https://doi.org/10.1111/j.1468-0084.2010.00591.x
37. Montiel Olea, J. L., & Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, *89*(4), 1789–1823. https://doi.org/10.3982/ECTA18756
38. Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708. https://doi.org/10.2307/1913610

39. Pástor, Ľ., Stambaugh, R. F., & Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, *142*(2), 550–571. https://doi.org/10.1016/j.jfineco.2020.12.011
40. Pástor, Ľ., Stambaugh, R. F., & Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, *146*(2), 403–424. https://doi.org/10.1016/j.jfineco.2022.07.007
41. Plagborg-Møller, M., & Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, *89*(2), 955–980. https://doi.org/10.3982/ECTA17813
42. Romano, J. P., & Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association*, *100*(469), 94–108. https://doi.org/10.1198/016214504000000539
43. Stock, J. H., & Watson, M. W. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics*, *20*(2), 147–162. https://doi.org/10.1198/073500102317351921
44. Zerbib, O. D. (2019). The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of Banking & Finance*, *98*, 39–60. https://doi.org/10.1016/j.jbankfin.2018.10.012
45. Zhao, D., Wang, Y., & Fang, Y. (2024). Greenium and public climate concerns: Evidence from China. *Finance Research Letters*, *69*., Article 106091. https://doi.org/10.1016/j.frl.2024.106091

Figures

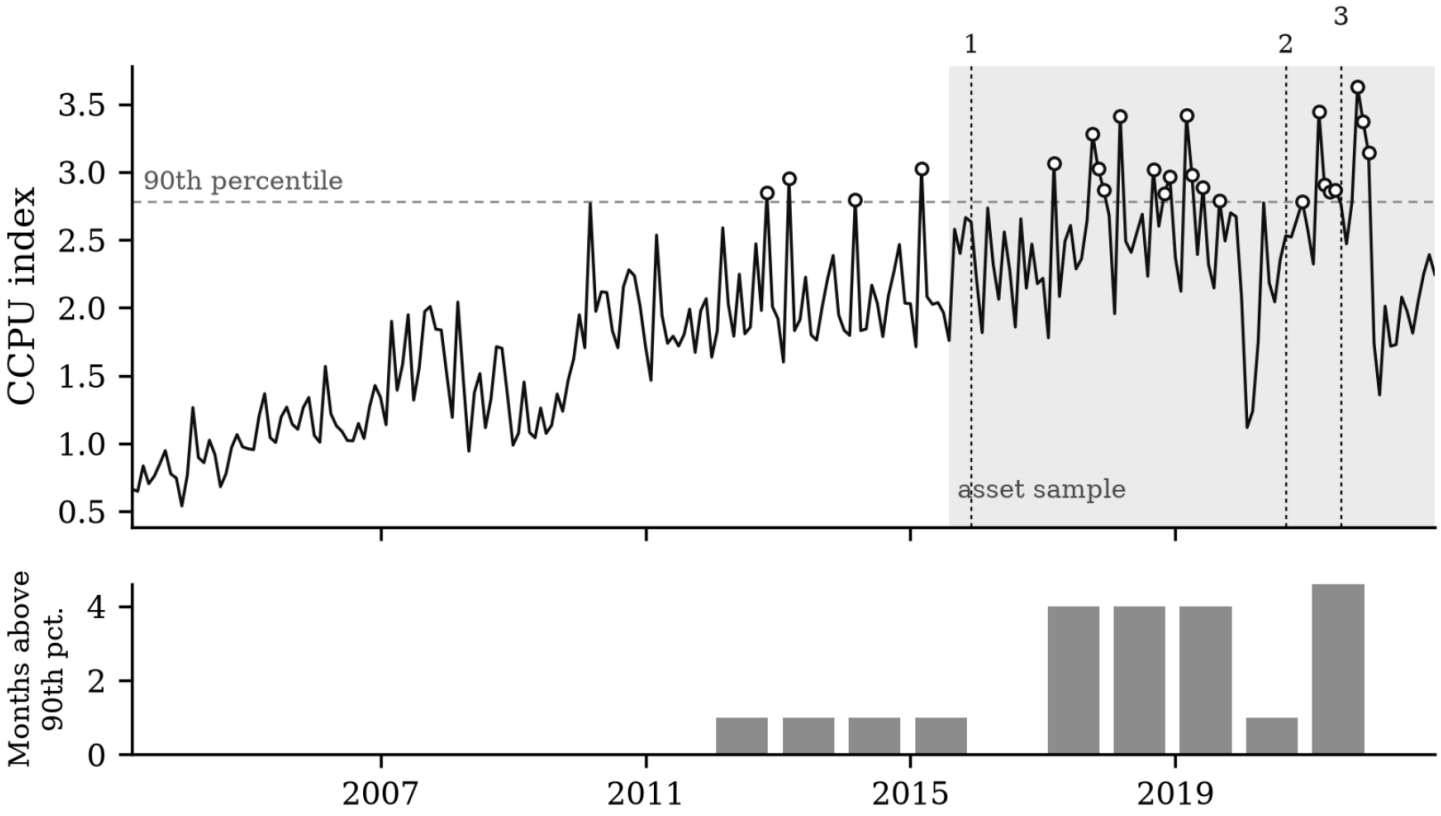


Figure 1

Climate policy uncertainty in China, 2003m4–2022m12. The dashed line is the ninetieth percentile of the full sample; open circles mark months at or above it. The shaded region is the asset-pricing sample. The lower panel counts the months at or above the ninetieth percentile by year. Vertical markers: (1) December 2015, launch of the supply-side reform; (2) September 2020, announcement of the dual-carbon target; (3) July 2021, opening of the national emissions trading scheme.

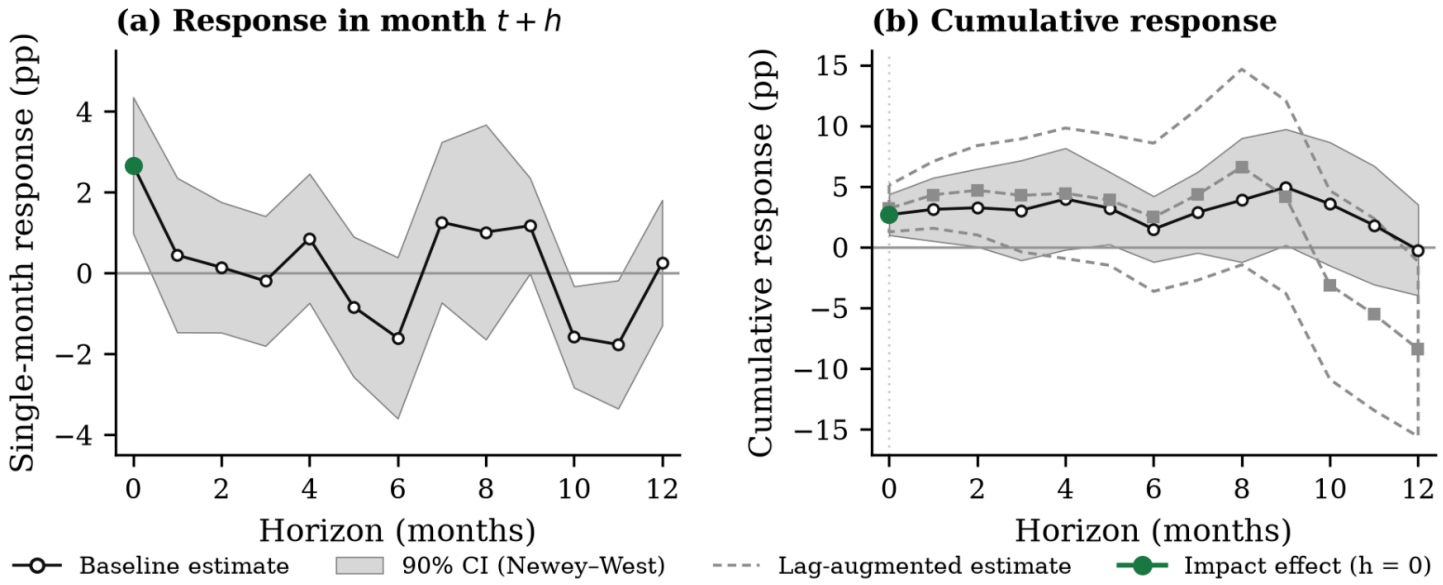


Figure 2

Response of the green–brown gap to a one-standard-deviation innovation in climate policy uncertainty. Panel (a) is the response in month t + h itself; panel (b) is the cumulative response. Shading in both panels is the 90 percent Newey–West interval. In panel (b) the dashed line and its outline are the lag-augmented local-projection estimate and its 90 percent interval. Filled markers are those with a t-statistic beyond 1.96.

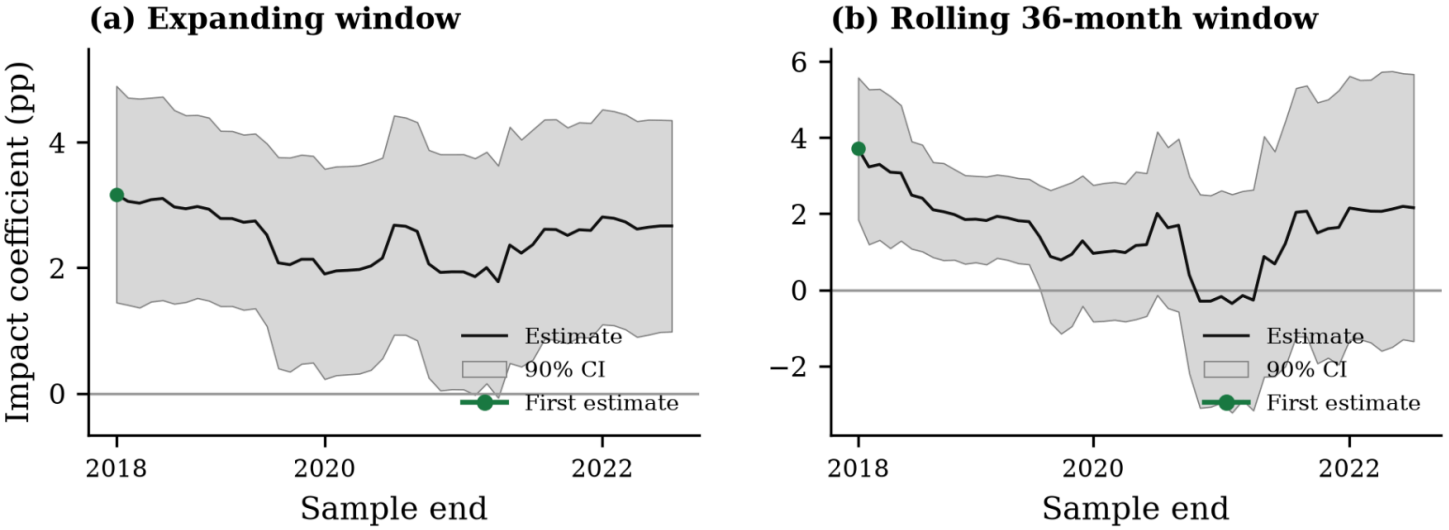


Figure 3

Stability of the impact coefficient. Panel (a): estimate on an expanding sample ending in the month shown. Panel (b): estimate on a rolling 36-month window ending in the month shown. Shading is the 90 percent Newey–West interval. The filled marker is the first estimate in each panel.

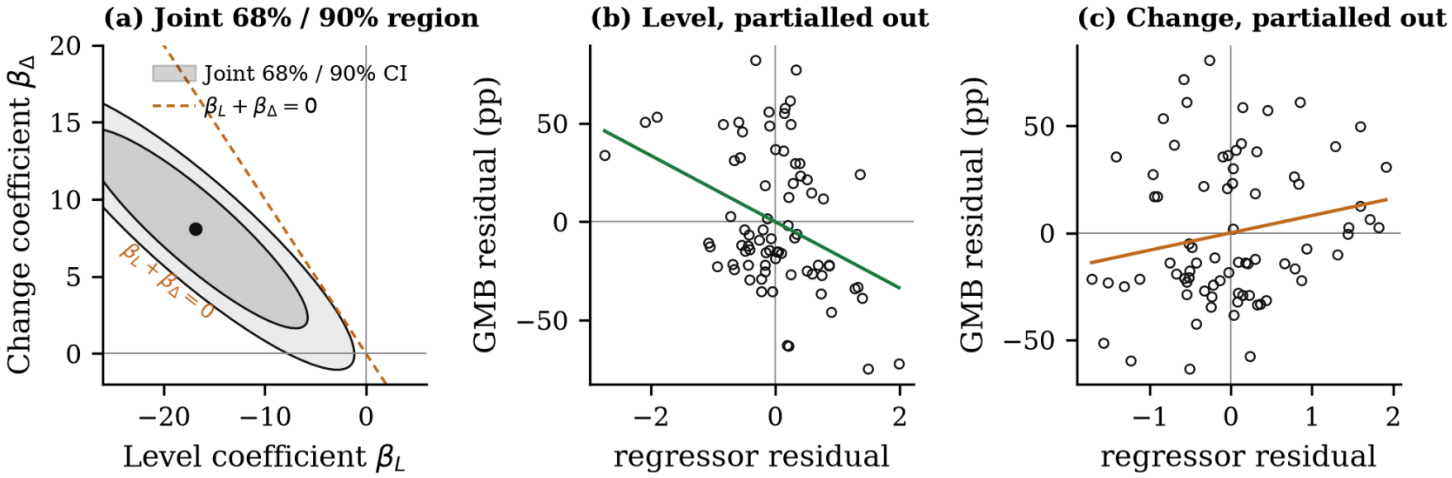


Figure 4

The twelve-month coefficient across ten specifications, grouped by the regressor whose coefficient is reported. Bars are 90 percent Newey–West intervals; filled markers are those with a t-statistic beyond 1.96 in absolute value, hollow markers those below it, a threshold that Section 5.5 shows to be too generous for the level at this horizon. All ten specifications use the same 76 observations at h = 12. The dashed horizontal line separates the seven specifications that report a coefficient on the level of the index from the three that report a coefficient on its change.

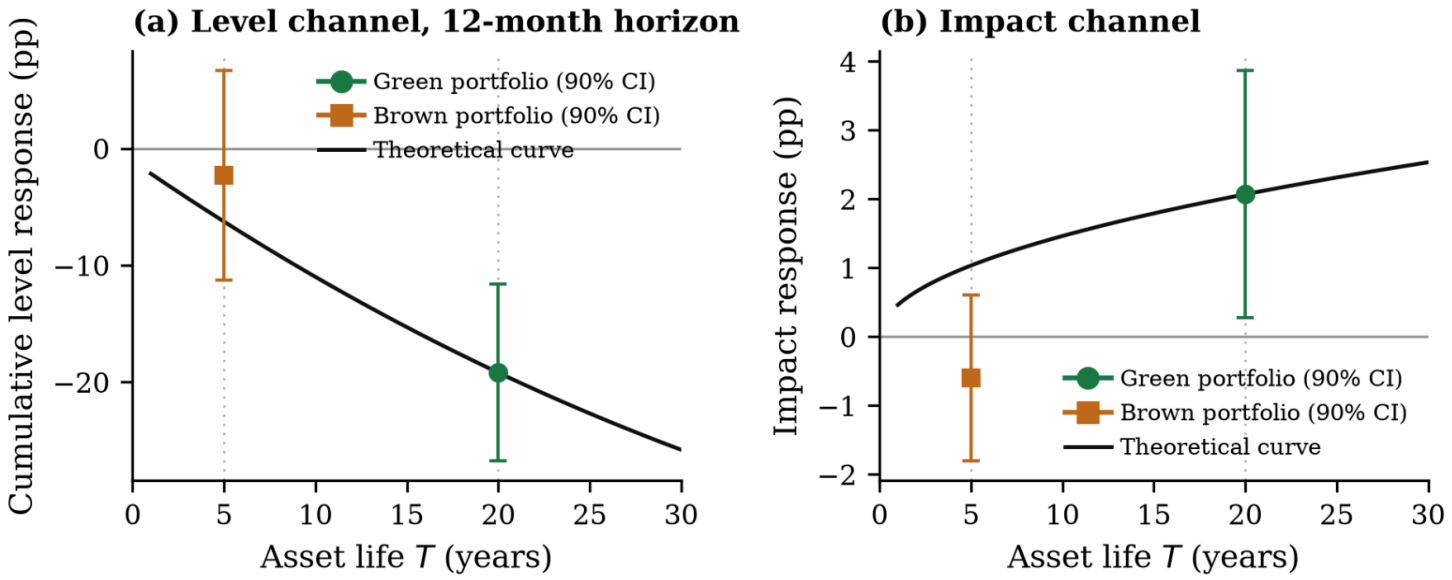


Figure 5

The level and change coefficients at h = 12. Panel (a): joint 68 and 90 percent confidence regions for the two coefficients, with the locus on which they are equal and opposite (dashed). Panels (b) and (c): added-variable plots, in which the vertical axis is the part of the twelve-month gap not explained by the other regressor and the horizontal axis is the part of each regressor not explained by the other; the plotted line is the corresponding coefficient.

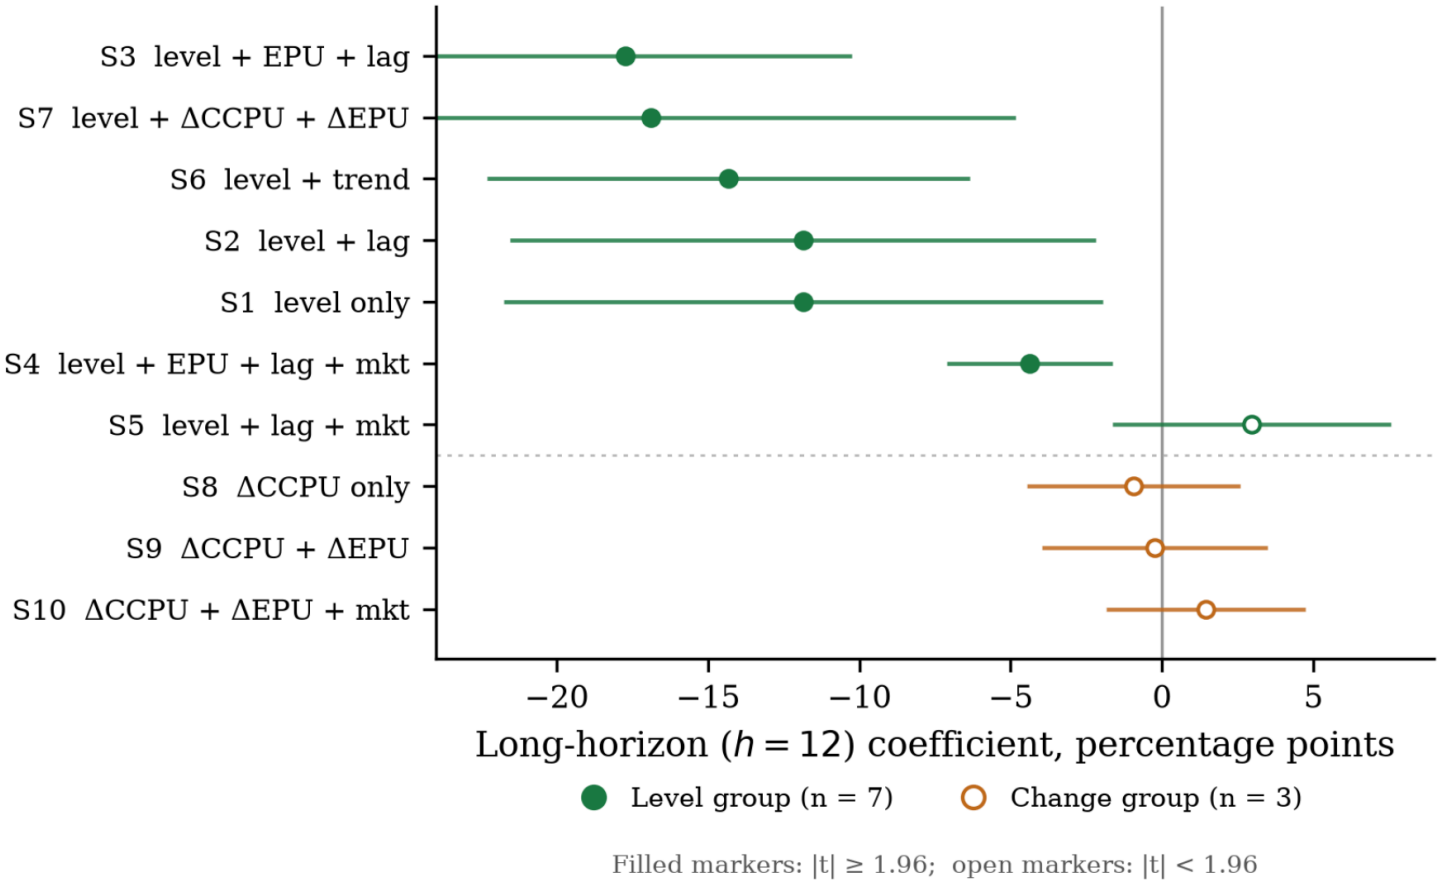


Figure 6

Two-sector comparative statics with the estimated points superimposed. Panel (a): semi-elasticity of the asset price to a persistent rise in the uncertainty premium, scaled so that the twenty-year asset matches the green estimate at h = 12. Panel (b): impact repricing, scaled so that the twenty-year asset matches the green estimate at h = 0. Filled circle: green portfolio; filled square: brown portfolio; bars are 90 percent intervals. Dashed vertical lines mark the illustrative green (20 years) and brown (5 years) horizons.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- tab4robust.csv

- tab3terms.csv
- tab2descriptives.csv
- tab5levelshock.csv
- tab6greenbrown.csv
- Appendix.docx