Discounting Chinese Industrial Policy*

Andrew J. Sinclair† Chuyi Zhang‡

September 10, 2026

Abstract

Each year, the Chinese central government issues dozens of policy documents directing its local officials to promote the development of targeted industries. Although central support is valuable in expectation, the duration and intensity of support are uncertain and revealed only over time. We find that small supported firms earn approximately 7.2% in annualized abnormal returns relative to a leading Chinese asset-pricing benchmark. We demonstrate that the pattern of these returns is consistent with rational discounting of policy-support uncertainty and is difficult to reconcile with alternative explanations based on systematic mispricing. Together, our results reveal an important market dynamic that is rooted in China's unique political economy.

JEL Classification: L52, P20, E02, G12

Keywords: Industrial policy, policy risk, risk premium, Chinese economy

*We thank Hanming Fang, Xiang Fang, Jerry Hoberg, Pawel Janas, Nan Jia, Wenxi Jiang, Raymond Kan, Toomas Laarits, Yang Liu, John Matsusaka, Jean-Laurent Rosenthal, Shang-Jin Wei, Luigi Zingales and seminar participants at the USC China Workshop, the Rotman School of Management, the HKU Business School, BlackRock, and the Hoover Institute for valuable advice and comments.

†California Institute of Technology, Baxter Hall, 1200 East California Blvd., Pasadena, CA 91125, United States. Email: ajsincla@caltech.edu.

‡Finance Area, Faculty of Business and Economics, The University of Hong Kong, Pokfulam Road, Hong Kong. Email: chuyi.zhang@connect.hku.hk.

1 Introduction

> The Party, the government, the military, civil society, and education,
>
> East, West, South, North, and the Center,
>
> **The Party leads everything**.[1]
>
> - Chinese Communist Party slogan

China is a vast country; it has a landmass comparable to the United States and a population more than three times the size of the European Union. But despite China's enormous size, the slogan above is perhaps much closer to reality than one might expect.

The Chinese central government "leads everything" because it is endowed with discretion over a vast amount of society's wealth. The state owns or controls all the land in China,[2] most of the banking sector (Sun, 2020), as well as over 320,000 state-owned enterprises. These resources are dispersed across the country, and are at the discretion of local government officials (Xu, 2011, 2025). To direct these resources, the central government issues policy documents to communicate to its local agents, as well as to inform the public, how resources are to be allocated (Cheung, 2014).

Between March 2008 and May 2022, the State Council, the de facto highest organ of state power in China, issued 282 such documents providing clear instructions to local officials to support the development of various target industries. These policies have been a core component of China's development model and, during the period under study, over one half of Chinese industries received policy support at one time or another.

[1] 党政军民学 东西南北中 党是领导一切的.

[2] Article 10 of the Constitution of the People's Republic of China states "Land in the cities is owned by the State. Land in the rural and suburban areas is owned by collectives except for those portions which belong to the State as prescribed by law; house sites and privately farmed plots of cropland and hilly land are also owned by collectives." Collectives are under the de facto control of the state.

Firms that receive policy support from the central government experience meaningful economic benefits. Supported firms pay less in taxes, gain improved access to short-term debt, report higher operating revenue, increase research and development spending, and exhibit less downside market risk. Market prices aggregate these various benefits and, on average, the stock price of a supported firm increases by approximately 150% over the five years following the initial release of a policy document.

However, these expected benefits also carry a high degree of uncertainty. Under China's discretion-based economic system, the central government has the authority to dynamically adjust and change policies over time. What this means is that, unless support is regularly reiterated, both local officials and the general population may believe that priorities have changed and support has quietly disappeared. Indeed, in over 60% of our cases, support disappears after just one year and stock returns level off. Sustained gains accrue only to firms that receive continued support.

In this paper, we explore Chinese policy risk from the perspective of stock market investors. To frame our analysis, we build a parsimonious asset-pricing model that captures the unique institutional features of China's economic system. Government policy support is valuable but uncertain, and this uncertainty may be resolved only gradually over time. Our model decomposes factor-adjusted returns into a policy-risk premium, revisions to beliefs about the implementation and continuation of support, and idiosyncratic news. Under rational expectations, belief revisions have a conditional mean of zero, while idiosyncratic news averages out in a sufficiently diversified portfolio. Provided that the benchmark model captures other priced risks, positive average abnormal returns can therefore be interpreted as compensation for policy-support risk.

Empirically, we find that small supported firms outperform standard asset-pricing bench-

marks and earn abnormal returns of approximately 7.2% per year under our preferred China-specific model. Consistent with a risk-premium interpretation, the policy portfolio performs poorly during high-marginal-utility states in China. Abnormal returns also vary with exposure to policy-support risk: returns are concentrated among private rather than state-owned firms, are stronger when central and local policy priorities align, and are present among small but not large firms. Together, these state-dependent and cross-sectional patterns are consistent with compensation for exposure to policy-support risk.

Systematic mispricing provides the principal alternative interpretation of our results. Because government policies are endogenous, the central government may select industries based on private information about their future prospects—or simply "pick winners." If investors fail to incorporate this signal fully, or if market frictions prevent them from acting on it, subsequent abnormal returns could reflect a gradual correction of mispricing rather than compensation for policy-support risk.

We investigate three broad mispricing channels: investor inattention, investor exuberance, and limits to arbitrage. In contrast to an inattention explanation, nearly 60% of policy announcements receive rapid coverage on China's national nightly news, and abnormal returns are statistically significant among covered rather than uncovered policies. Industry news coverage and trading activity are also elevated around policy announcements. Regarding the exuberance explanation, margin-financed trading is not unusually elevated around around policy announcements, and we find no evidence of the long-term return reversals that would follow if investors temporarily bid prices up too high. Finally, several prominent limits to arbitrage appear unlikely to sustain the return pattern: the strategy is long-only, requires no leverage, and survives tests designed to limit the influence of illiquid stocks.

We also examine whether the estimated abnormal returns reflect insufficient diversifica-

tion across supported industries. In particular, the strategy may load on industries that happened to perform unusually well during the sample period. We therefore reconstruct the portfolio after separately excluding the technology sector and the top decile of industries ranked by realized performance. The abnormal return remains positive and statistically significant in both tests. Our most stringent test—excluding the top-performing decile—produces a conservative estimate of approximately 3.6% annualized.

We next examine whether the abnormal returns reflect risks omitted from our benchmark asset-pricing models. Our preferred specification uses the leading China-specific asset-pricing model (Liu et al., 2019), which we augment with factor-mimicking portfolios for aggregate economic policy uncertainty (EPU) (Baker et al., 2016; Brogaard and Detzel, 2015) and R&D investment risk (Leung et al., 2020). The estimated abnormal return remains economically and statistically significant after controlling for each additional factor.

Finally, we examine whether the results are sensitive to how policy support is measured. Our baseline measure uses a transparent title-keyword rule to classify supportive documents and maps official policy themes to industries. Alternative measures classify the full text using manual reading, ChatGPT, and Chinese-language DeepSeek models. Although the resulting classifications differ, each produces significant abnormal returns among small supported stocks. The DeepSeek analysis further shows that the returns are concentrated among policies promising concrete resource support rather than vague guidance.

Our paper makes three core contributions. First, we contribute to the literature on political uncertainty and asset prices. Pástor and Veronesi (2012, 2013) show theoretically that uncertainty about government policy can affect asset prices and command a risk premium, while Kelly et al. (2016) provide empirical evidence from equity options surrounding elections and global political summits. We study a different form of political uncertainty arising

from ongoing discretionary government action, where uncertainty is resolved gradually rather than around a discrete political event. Our findings suggest that the Pástor–Veronesi mechanism extends beyond discrete political events: policy uncertainty can be priced even when government action unfolds gradually and its resolution has no predetermined date.

Second, we contribute to the growing literature on the political economy of Chinese finance (Brunnermeier et al., 2017, 2022; Sinclair, 2026). Brunnermeier et al. (2022) show that intensive state intervention can induce investors to produce information about government actions rather than economic fundamentals. We emphasize a complementary implication of China's discretion-based economic system: institutions also shape the risks investors bear. When firm value depends on discretionary government support, uncertainty about how that discretion will be exercised becomes relevant for expected returns.

Third, we provide the first systematic analysis of industrial-policy risk. The industrial-policy literature has largely focused on the effects of particular policies or policy episodes (Aghion et al., 2015; Barwick et al., 2019; Kalouptsidi, 2018; Liu, 2019). We show that industrial policy both provides meaningful benefits and creates industry-level risk because the value of these benefits depends on the implementation and continuation of support. Our asset-pricing methodology allows us to analyze a wide range of heterogeneous policies by aggregating the effects of policy support on firm valuations. In doing so, we provide large-scale empirical evidence on an important consequence of industrial policy that is difficult to capture through the analysis of individual policy episodes.

The remainder of this paper is organized as follows. Section 2 discusses the economics of Chinese policy and presents a parsimonious model of policy risk. Section 3 describes the construction of our database. Section 4 documents the value and uncertainty of policy support. Section 5 estimates abnormal returns to supported firms. Section 6 evaluates

whether these returns are consistent with compensation for policy-support risk. Section 7 evaluates alternative interpretations and presents additional robustness tests. Section 8 concludes.

2 Institutional Setting and Conceptual Framework

2.1 The Institutions of Chinese Policy

For our analysis, it will be helpful to consider industrial policy from a comparative institutional perspective. Under the institutions of the People's Republic of China, the central government is endowed with a high degree of authority to act. This contrasts with the institutions of Western liberalism, in which the rule of law places significant constraints on government actions. This distinction between China's discretion-based system and the West's rules-based system has important implications in a dynamic setting, as emphasized by the comparative rules-versus-discretion literature (Brunnermeier et al., 2017; Kydland and Prescott, 1977).

In particular, under a rules-based system, rules do not need to be reiterated every period. Society understands that a rule that was enacted in the past still holds today; for instance, it is clear to every American that the government cannot arbitrarily abolish the Second Amendment tomorrow. Under such a system, rules are persistent and often constrain the ability of the government to intervene.

Under a discretion-based system, rules on basic social conduct, such as prohibitions against murder and theft, do not need to be reiterated. But government policies on discretionary issues, such as social planning (e.g., the One Child Policy), party discipline (e.g., the Anti-Corruption Campaign), and industrial policy, need to be frequently reiterated. If these

policies are not reiterated, both local officials and the public may believe these policies are no longer top priorities and enforcement or implementation may wane.

Policy documents, such as the State Council policies studied in this paper, are thus a central feature of how the Chinese system works. As the opening quote to this paper suggests, these policies are part of the economics of centralized leadership and steer all aspects of society. Just as laws are the foundation of a rules-based governance system, policies are the foundation of a discretion-based governance system.

2.2 Gradual Information Revelation

For our analysis, the difference between the permanence of rules and the flexibility of discretionary policy is paramount. In a rules-based system, investors understand that industrial policy will persist for the stated lifetime of that policy. In a discretion-based system, a policy announced today may change tomorrow.

In particular, the flexibility to dynamically adjust policies is a core part of China's development model. This approach is epitomized by the maxim "crossing the river by feeling the stones" (摸着石头过河): policies can first be piloted, adjusted in light of actual conditions on the ground, and then expanded or contracted based on results.

This flexibility, however, creates uncertainty for investors. China's electric-vehicle industry provides an important example. Although the development of this industry has been fueled by industrial policy support, investors were uncertain about the future path of policy and wary of committing capital. In 2016, the head of manufacturing at NextEV explained his company's reticence: "We don't really want to spend too much money on manufacturing, to be honest, and the government policy might change."[3] Under a rules-based system, the

[3]Reuters, "China Toughens Up on Producing Fewer, Better Electric Car Makers," November 18, 2016,

government can credibly commit to providing support in the future; but in a discretion-based system, no such commitment is possible.

When investing in a supported industry, investors must thus ask the question: will policy support continue? Over time, if the government continues to promote an industry, investors will gradually revise beliefs upward. However, if support is not reiterated, investors and the public will gradually infer that the industry is no longer a government priority.[4]

2.3 Conceptual Framework

We formalize this logic using a parsimonious model that captures how policy-support uncertainty affects expected returns.

Let the terminal payoff of firm i be

$$X_{i,T} = F_{i,T} + B_i S_{i,T}, \tag{1}$$

where $F_{i,T}$ is the payoff absent effective policy support, $B_i > 0$ is the net value of the policy benefit, and $S_{i,T} \in \{0, 1\}$ indicates whether that benefit remains effective at horizon T.

For simplicity, we represent the stream of policy benefits by a single terminal payoff. We use a binary variable to isolate uncertainty about implementation and continuation. More generally, $B_i S_{i,T}$ could be modeled as a random net policy payoff that is positive in expectation but can take negative values.[5]

available via *Fortune*.

[4]Appendix E.1 documents that the State Council rarely formally withdraws support for an industry.

[5]Negative realizations would capture reciprocal obligations or the possibility that government support comes with "strings attached."

At time $t < T$, investors assign probability

$$p_{i,t} = \Pr_t(S_{i,T} = 1) \tag{2}$$

that the policy benefit will be implemented and remain effective at horizon T. News about policy implementation and continuation produces the revision

$$\Delta p_{i,t+1} = p_{i,t+1} - p_{i,t}. \tag{3}$$

Under rational expectations, posterior beliefs satisfy

$$E_t[\Delta p_{i,t+1}] = 0. \tag{4}$$

Let m_{t+1} denote the stochastic discount factor and $P_{i,t}$ the firm's price. The policy-related expected-return component can thus be written as

$$\mu^P_{i,t} = -\frac{B_i}{P_{i,t}} \frac{\mathrm{Cov}_t(m_{t+1}, \Delta p_{i,t+1})}{E_t[m_{t+1}]}. \tag{5}$$

A positive premium arises when adverse revisions about the effectiveness or continuation of support occur in states in which marginal utility is high. In that case, $\Delta p_{i,t+1}$ is low when m_{t+1} is high, making the covariance negative.[6]

To connect the model to our empirical tests, define the factor-adjusted excess return as

$$\widetilde{R}_{i,t+1} = R_{i,t+1} - R_{f,t+1} - \boldsymbol{\beta}_i' \mathbf{f}_{t+1}, \tag{6}$$

[6]Furthermore, policy-support uncertainty need not increase conventional volatility: support can protect firms from ordinary downside outcomes even while the loss of that protection remains especially costly in particular aggregate states.

where $\mathbf{f}_{t+1}$ contains the benchmark factor returns. We write this return as

$$\widetilde{R}_{i,t+1} = \underbrace{\mu_{i,t}^{P}}_{\text{policy-risk premium}} + \underbrace{\frac{B_i}{P_{i,t}}\Delta p_{i,t+1}}_{\text{policy news}} + \underbrace{\varepsilon_{i,t+1}}_{\text{other news}}, \tag{7}$$

where $\varepsilon_{i,t+1}$ captures idiosyncratic innovations to fundamentals and other return variation.

The belief revision $\Delta p_{i,t+1}$ captures both gradual information revelation and revisions associated with policy announcements. Policy documents need not constitute large, discontinuous news events. In China, a policy document often marks one stage of a longer deliberative and implementation process involving both the government and private-sector participants. The contents of a newly released document may therefore have been anticipated from earlier speeches, development plans, or consultations.[7]

Taking conditional expectations gives

$$E_t\left[\widetilde{R}_{i,t+1}\right] = \mu_{i,t}^{P} + \frac{B_i}{P_{i,t}} E_t\left[\Delta p_{i,t+1}\right] + E_t\left[\varepsilon_{i,t+1}\right] \tag{8}$$

Under rational beliefs, policy-related belief revisions have a conditional mean of zero. This means that endogenous government selection does not mechanically produce predictable post-announcement returns under rational expectations. If investors understand the government's selection rule, the information contained in that selection should already be reflected in prices.[8] Idiosyncratic news also has a mean of zero and averages out in a sufficiently diver-

[7] Consistent with this, policy announcements are closely followed in the media (see Section 7.1.1) but nevertheless have limited immediate price impact (see the event study in Appendix F).

[8] Section 7.1 examines the possibility that investors do not fully incorporate this information.

sified portfolio. The expected factor-adjusted return then equals the policy-risk premium:

$$E_t\left[\widetilde{R}_{i,t+1}\right] = \mu_{i,t}^{P}. \tag{9}$$

Our main analysis in Section 5 estimates average factor-adjusted returns and shows that the return to a portfolio of small supported firms is positive and significant. Section 6 examines whether this return behaves like compensation for policy-support risk. The portfolio performs poorly in high-marginal-utility states, consistent with $\mathrm{Cov}_t(m_{t+1}, \Delta p_{i,t+1}) < 0$, and abnormal returns vary with exposure to policy risk. Section 7 then evaluates the core assumptions behind this interpretation. Systematic errors in investor beliefs could produce a nonzero average realization of policy news, omitted risk factors could generate abnormal returns relative to our benchmarks, and concentrated industry outcomes could prevent idiosyncratic news from averaging out. We examine each possibility empirically.

3 Data and Measurement

3.1 Policy-Document Database and Industry Mapping

We obtain policy records from the official State Council policy-document database.[9] The database contains documents from the State Council and ministries and commissions under the State Council; our empirical sample focuses on documents issued by the State Council or the General Office of the State Council. For each record, we collect the title, full text, issuing agency, publication date, origination date, and official policy theme. The initial collection contains 5,762 records published between October 1996 and July 2022.

[9] `https://sousuo.www.gov.cn/zcwjk/policyDocumentLibrary`

Publication dates before March 28, 2008 are unreliable: 969 older records share that publication date, and the average gap between origination and publication is substantially longer before that point. We therefore begin with the 2,663 documents published after March 28, 2008.[10] We identify substantive policy instruments using administrative terms in document titles, including *decision*, *measure*, *opinion*, and *plan*. This filter removes routine notices, approvals of matters without a clear economic-policy content, and records invalidating earlier documents, leaving 1,574 policy documents. Appendix A reports the complete filter.

The State Council database assigns documents to 95 official themes (see Panel A of Table 1). We manually map themes to industries using the CSRC's 2012 Industry Classification of Listed Companies. Thirty-four themes can be mapped to at least one industry; some themes, including technology and intellectual property, map to several industries. Many unmapped themes concern administrative or social matters, such as population, religion, cultural affairs, public security, or Taiwan affairs, and cannot be systematically mapped to an economic industry.[11] Panel B of Table 1 reports the theme-to-industry mapping.

This procedure yields 505 mapped policy documents. Manual review removes 47 records that are not substantive policies and 54 documents directed at individual regions, leaving 404 policies published between April 24, 2008, and May 27, 2022.

3.2 Measuring Supportive Policy

Our baseline classifier identifies support using favorable language in document titles. Terms such as *support*, *promote*, *develop*, *encourage*, and *accelerate* indicate that a document directs favorable attention toward the mapped industry. Applying the title rule to the 404-policy

[10] Naughton (2021) also dates the emergence of China's modern industrial-policy system to the period beginning in 2006.

[11] The mapping therefore focuses the analysis on policies with an observable listed-firm exposure and does not imply that policies under unmapped themes lack economic importance.

sample identifies 282 supportive documents. Appendix A provides the complete word list and methodology.

The baseline classification rule is transparent and does not use information about portfolio returns, but may not fully reflect the content of the documents or the intensity of support. We therefore use three additional full-text classifications to assess and refine the baseline measure. The first is a manual full-text classification. The second is a ChatGPT-based classification that reads the body of each policy document and evaluates whether it provides support for the mapped industries. The third is a Chinese-language extraction based on DeepSeek. Most policy–industry links classified as supportive by the AI-based measures are also supportive under manual coding: 80.2% for ChatGPT and 79.4% for DeepSeek V4-Pro. For details on these alternative classification schemes, see Appendix A.

Our DeepSeek algorithm also extracts content from the policy documents (see Appendix A.4). We find that 241 of the 404 policies (59.7%) provide resource-oriented support, including fiscal transfers, credit, tax provisions, public procurement, pricing, market access, mandated demand, and resource allocation; 31 policies (7.7%) provide executable nonresource support, such as regulatory and administrative actions; finally, 38 policies (9.4%) provide only vague support. The remaining 94 policies do not contain favorable support under this classification.[12]

3.3 Financial Data and Sample Construction

We obtain stock returns, trading activity, and firm financial data from CSMAR. The sample includes A-share firms listed on the main boards of the Shanghai and Shenzhen Stock

[12]In Panel C of Table 9, we find that positive and statistically significant abnormal returns are concentrated among policies that direct resource support; policies providing only vague guidance do not produce significant abnormal returns.

Exchanges and on the Growth Enterprise Market. Following Liu et al. (2019), we exclude firms during the first six months after listing and require at least 120 trading days during the preceding 12 months and at least 15 trading days during the preceding month.[13] The underlying firm sample spans 2008 through 2021 and contains 4,288 firms across 81 industries. During this period, 2,908 firms in 53 industries receive policy support.

The annual firm-outcome analysis uses CSMAR balance-sheet and income-statement variables. We exclude financial firms and winsorize annual variables at the 1st and 99th percentiles. Section 4.1 defines the outcome-specific controls and fixed effects. Section 5.1 describes monthly portfolio formation, the size split, return weighting, and the asset-pricing factors.

3.4 Sample Coverage and Summary Statistics

Figure 1 reports the distribution of the 282 supportive policy documents by year and shows a pronounced increase in 2015 and 2016, coincident with the launch of the 13th Five-Year Plan. Panel A of Figure 2 reports that the number of supported industries during the sample ranges from 5 to 42, with an average of 22 out of 81 industries. Panel B shows that the number of listed firms in industries supported during the preceding 12 months ranges from 99 to 1,718, with an average of 967. The few large discontinuities reflect broad technology policies that map to many industries; Section 7.2 shows that the main results survive the exclusion of these policies.

Table 2 presents annual firm-level summary statistics after winsorization. The median firm is slightly less than eight years old and has assets of approximately RMB 2.6 billion.

[13]We do not exclude ST and ST* firms. We investigated these stocks and found that ST and ST* firms account for less than 2% of the observations. Removing these firms leaves the results almost unchanged, both qualitatively and quantitatively.

Annual stock returns have a median of -0.4%, a mean of 13.8%, and a cross-sectional standard deviation of 59.0%. In the firm-month sample, 37.7% of observations are classified as supported based on at least one supportive document during the preceding 12 months. The policy-support measure therefore covers a broad but time-varying cross-section of listed firms.

4 The Economic Value and Uncertainty of Policy Support

State Council policy documents direct local officials and state-linked institutions to support targeted industries, but the economic consequences of these directives are not guaranteed to persist. In this section, we first show that policy support is associated with meaningful changes in firms' access to resources, operating activity, investment, and downside risk. We then document that support is uncertain: support is frequently not renewed after the first year but becomes more predictable as it persists. Finally, we demonstrate that the returns to supported firms correlate with the duration of support, suggesting that policy-support uncertainty may affect investor returns.

4.1 Policy Support and Firm-Level Outcomes

We begin by testing whether the policy documents in our sample are associated with economically meaningful changes at the firm level.[14] For each firm-year, we measure the strength of policy support, $SupportStr_{i,t}$, as the number of supportive policies applicable to firm i's industry divided by the total number of supportive policies issued in year t. The unconditional mean of this measure is 4.5%, with a standard deviation of 7.0 percentage points. We

[14]The full analysis and tables are presented in Appendix C.

estimate the following fixed-effects specification:

$$Y_{i,t+1} = \beta_1 SupportStr_{i,t} + \boldsymbol{\beta}_2' \mathbf{X}_{i,t} + a_i + a_t + u_{i,t}, \tag{10}$$

where $Y_{i,t+1}$ measures firm i's financial or operating outcome in the following year. Depending on the outcome, the control vector $\mathbf{X}_{i,t}$ includes firm size, age, leverage, asset tangibility, political connections, operating revenue, and return on assets. All specifications include firm and year fixed effects, and standard errors are clustered by industry. We exclude financial firms from this analysis.

Policy support is associated with greater access to resources. Panel A of Table 3 shows that a one-standard-deviation increase in support strength is associated with a 0.36 percentage point decline in the effective tax rate,[15] a 5.07 percentage point increase in cash growth, and a 5.73 percentage point increase in short-term borrowing growth. By contrast, long-term borrowing does not increase significantly. This pattern is consistent with central policy directives improving firms' access to short-term credit and liquidity through local governments and state-linked financial institutions.

Policy support also predicts changes in firms' operations and investment. Panel B of Table 3 shows that stronger support is associated with higher operating-revenue growth. Supported firms also expand R&D spending: a one-standard-deviation increase in support strength is associated with a 3.35 percentage point increase in R&D spending growth. These findings indicate that policy support is followed by greater operating activity and increased investment in innovation.

These changes do not translate mechanically into higher profits. Operating profit and net

[15] To obtain these statistics, multiply a one-standard-deviation change in support strength (0.07, from Table 2) by the point estimate. For example, the estimated change in the effective tax rate is $0.07 \times (-0.051) = -0.00357$, or -0.36 percentage points.

income decline following stronger policy support, and selling expenses rise, while financial expenses do not change significantly. Greater borrowing, revenue, and R&D spending therefore indicate that policy support is associated with changes in firms' resources and behavior, but these outcomes alone do not establish an unambiguous increase in shareholder value.

The evidence is stronger, however, when policies provide concrete support. In Appendix D, we analyze the firm-level outcomes associated with resource-oriented support. Resource-oriented support is associated with lower effective tax rates, greater cash growth, higher short-term borrowing, and *higher* subsequent operating profits. These patterns are consistent with supportive policies translating into tangible firm-level benefits through fiscal and financial channels.

Policy support may also benefit targeted firms by limiting adverse outcomes. We therefore supplement the accounting outcomes with measures of downside risk and test whether downside risk changes following supportive-policy announcements.

The results are presented in Table 4. The measures capture several distinct dimensions of downside exposure: the severity of adverse outcomes (maximum drawdown and the 5th-percentile loss), the frequency and dispersion of negative returns (the incidence of daily losses exceeding 5 percent and lower-partial volatility), and measures of crash risk (NCSKEW and DUVOL). Across all six measures, downside risk is lower in the post-announcement period. For the first announcement (Panel A), the reductions are economically sizable and statistically significant for five of the six measures. The larger sample of all supportive-policy announcements (Panel B) produces the same pattern with substantially greater precision: all six measures decline significantly after the announcement. Consistent with a protective role for government support, these results show a broad post-announcement reduction in the severity, frequency, and asymmetry of downside market risk.

4.2 Continuation and Duration of Policy Support

The value of policy support depends not only on the content of a directive, but also on how long the central government maintains the underlying priority. A single document can initiate access to credit, tax relief, administrative assistance, or other resources, but many firm responses require support to remain in place. At the time of initial support, firms and investors can observe the current policy direction but cannot observe the eventual duration of that support.

Panel A of Figure 3 documents substantial uncertainty at the beginning of a support episode. An industry receiving support for the first time has approximately a 40% probability of receiving support again in the following year. Thus, most initial episodes are not followed by a supportive policy document in the next year. Conditional continuation probabilities rise sharply once support persists: after two consecutive years of support, the probability of a third year is approximately 75%, and after seven consecutive years, continued support becomes nearly certain. The central government's commitment to an industry therefore becomes more predictable as repeated directives accumulate, while continuation is most uncertain near the beginning of a support episode.

The realized duration distribution reinforces this pattern. Panel B of Figure 3 reports the maximum uninterrupted support duration for each industry. More than one-third of industries receive support for only one year, while nearly one-quarter receive support for eight years or more. The distribution is therefore sharply bimodal: support often either disappears quietly or develops into a persistent central priority. At initial issuance, however, the eventual location of an industry within this distribution is not yet known.[16]

[16] Appendix B demonstrates that this bimodality also exists for industries that are included in the national Five-Year Plans. That is, even top national priorities face policy-support uncertainty.

These duration patterns establish an important source of policy-support uncertainty. Section 4.1 showed that policy support is associated with tangible firm-level outcomes, but this section demonstrates that investors face uncertainty about whether support will be reiterated in the future. Over time, uncertainty about continuation declines. Investors observe whether support accumulates and p approaches one, or whether support quietly disappears and p approaches zero.

4.3 Realized Value and Continued Support

We next examine whether support continuation corresponds to the realization of firm-level and market outcomes. For the firm-outcome and downside-risk comparisons in Appendix E.2, we identify the first policy announcement for an industry and define continued industries as those that receive support in the following year, while stopped industries receive no support during the following two years.

Appendix E.2 compares firm-level outcomes for firms with stopped versus continued policy support. The comparisons in Table E2 show lower operating-revenue and operating-profit growth among firms whose support stops. These differences suggest that support renewal is associated with economically meaningful differences in subsequent operating performance.

We also find evidence that continued support provides more downside-risk protection. Table E3 shows that firms in the continued-support group experience greater improvements in several downside-risk measures than firms in the stopped-support group. This pattern is consistent with continued support providing greater downside-risk protection.

We next demonstrate that future stock returns are correlated with the duration of support. Panel A of Figure 4 presents the average returns following initial policy support. On average, stock returns increase by over 150% in the five years following initial support. How-

ever, as Panel B demonstrates, returns diverge sharply based on the duration of support. Returns are similar during the first year but subsequently diverge. Panel C statistically tests this relationship and demonstrates that returns are statistically indistinguishable in the first year, but that returns for the continuously supported group are statistically higher than the discontinued group in years two, three, and four.

Panel D examines average returns following the last supportive policy document and demonstrates that returns are still elevated in the first year, consistent with investors initially assigning a positive probability to continued support, but returns decline below the market average in subsequent years as it becomes increasingly clear that support has quietly ended ($p \to 0$).

Taken together, these results suggest that, on average, policy support is associated with benefits at the firm level, but the duration of such benefits is highly uncertain at the outset. The return analysis further shows that stock returns vary with the realized continuation of support.

5 Stock-Market Pricing of Policy Support

We next examine whether the returns of supported firms can be explained by standard risk factors in the Chinese market. We first describe the portfolio construction and asset-pricing benchmarks, and then document our central return result: small firms in supported industries earn positive abnormal returns relative to standard Chinese asset-pricing factors. This result implies that the sample average of the factor-adjusted return in equation 8 is positive for a diversified portfolio of supported firms.

5.1 Portfolio Construction and Asset-Pricing Benchmarks

At the beginning of each month, we identify firms in industries that received at least one supportive State Council policy document during the preceding 12 months. We form equal- and value-weighted portfolios of these firms and hold each portfolio for one month. The portfolios are rebalanced monthly as new policy documents enter the classification window and older documents leave it. Value weighting limits the influence of the smallest firms, while equal weighting captures the performance of the average supported stock.

We conduct the analysis separately for firms above and below the median monthly market capitalization among eligible firms. This distinction allows the relation between policy support and returns to vary with firm size. Chinese industrial policy is often meant to spur competition and thus may disproportionately benefit small firms.

We evaluate portfolio returns using three asset-pricing benchmarks. The Capital Asset Pricing Model controls for market exposure. The Carhart (1997) four-factor model adds size, value, and momentum factors. Our primary benchmark is the CH4 model of Liu et al. (2019), which was developed for the Chinese equity market and explicitly accounts for the unusually strong returns of small Chinese firms. We obtain the risk-free rate and the market, size, value, and momentum factors from CSMAR, and obtain the CH4 factors from Robert Stambaugh's website.[17]

For each portfolio p, we estimate

$$R_{p,t} - R_{f,t} = \alpha_p + \boldsymbol{\beta}_p' \mathbf{f}_t + \epsilon_{p,t}, \tag{11}$$

where $R_{p,t} - R_{f,t}$ is the portfolio's monthly excess return, $\mathbf{f}_t$ contains the factors included

[17] `https://finance.wharton.upenn.edu/~stambaug/`

in the relevant benchmark, and $\epsilon_{p,t}$ is the regression residual. The intercept, α_p, measures the average portfolio return not explained by those factors. A positive alpha indicates an average abnormal return relative to the specified benchmark.

The estimated alpha therefore measures the sample mean of the factor-adjusted return defined in equation 8.

5.2 Abnormal Returns to Supported Firms

Table 5 reports abnormal returns to portfolios of supported firms.[18] The estimates are positive and statistically significant for small firms across all three asset-pricing benchmarks and under both weighting methods. In the value-weighted portfolio, monthly alpha is 1.0% under the CAPM, 0.5% under the Carhart model, and 0.6% under the CH4 model. The corresponding annualized estimates are 12.0%, 6.0%, and 7.2%, respectively.[19] The equal-weighted portfolio produces a similar CH4 alpha of 0.6% per month. Our preferred estimate is the CH4 alpha of approximately 0.6% per month, or 7.2% per year, because the CH4 model was designed specifically to account for the strong size and value effects in the Chinese equity market.

The abnormal returns are concentrated among small firms. For large supported firms, the equal-weighted CH4 model produces a positive monthly alpha of 0.5%, but the CAPM and Carhart estimates are not statistically significant. None of the value-weighted large-firm alphas are statistically significant. The contrast across firm size indicates that the headline return pattern is not a uniform feature of all firms in supported industries. Table 5 also reports a placebo portfolio of small non-supported firms. Non-support is the complement of support: it includes both firms exposed only to non-supportive policy documents and firms

[18]The full factor loadings are reported in Tables J1, J2, and J3 in Appendix J.

[19]We annualize by multiplying monthly alpha by 12.

exposed to no policy document during the preceding 12 months. This portfolio tests whether the headline result reflects the general small-firm anomaly in China.

Consistent with the well-documented small-firm anomaly in China, non-supported small firms earn positive CAPM and Carhart alphas. Conventional models have difficulty explaining the returns to small Chinese stocks, and Liu et al. (2019) developed the CH4 model specifically to address this issue. This is why the CH4 model is our preferred benchmark. Under the CH4 model, the small-firm anomaly is absorbed: the portfolio of small non-supported firms produces an alpha close to zero that is statistically insignificant under both weighting schemes. The final portfolio is a long-short portfolio of small supported-minus-non-supported firms. This portfolio earns positive alpha under all three benchmarks. The spread is positive even under the CAPM and Carhart models, in which the general small-firm anomaly has not been fully absorbed, and is largest under CH4.

These estimates establish the paper's central return fact: small firms in supported industries earn returns that standard Chinese asset-pricing factors do not explain. The estimated alpha measures the average portfolio return not explained by the benchmark factors (equation 8). Under rational beliefs, and provided that the benchmarks capture other priced risks, this alpha can be interpreted as compensation for policy-support risk. The remaining analysis examines the state dependence and cross-sectional structure of the returns and evaluates alternative interpretations.

6 Is Policy-Support Uncertainty a Priced Systematic Risk?

Section 5 shows that a diversified portfolio of small supported firms earns positive abnormal returns. Under rational beliefs, and provided that the benchmark models capture other

priced risks, these returns can be interpreted as compensation for policy-support risk. We now examine whether the state dependence and cross-sectional structure of the returns are consistent with this interpretation.

Recall from equation 5 that a positive policy-risk premium arises when revisions in beliefs about policy support are negatively correlated with the stochastic discount factor. That is, if $\mathrm{Cov}_t(m_{t+1}, \Delta p_{i,t+1}) < 0$, then the risk premium is greater than zero:

$$\mu_{i,t}^{P} = -\frac{B_i}{P_{i,t}} \frac{\mathrm{Cov}_t(m_{t+1}, \Delta p_{i,t+1})}{E_t[m_{t+1}]} > 0.$$

This expression produces two empirical predictions. First, the policy portfolio should perform poorly in high-marginal-utility states. Second, firms with greater exposure to policy-support uncertainty should earn higher abnormal returns. We examine these predictions in turn.

6.1 Policy Returns in High-Marginal-Utility States

We begin by asking: what is a high-marginal-utility state in China? Conventional asset-pricing tests often use market declines, high volatility, weak demand, and low household confidence as proxies for high marginal utility. These measures capture periods in which household consumption or wealth is likely to be especially valuable.

In China, however, the relation between conventional bad-state measures and marginal utility may be ambiguous. Under China's discretion-based economic system, the central government retains broad authority to intervene. A decline in the stock market or a deterioration in economic activity may worsen household conditions, but it may also increase expectations of government support. Because this expected intervention provides implicit insurance (Sinclair, 2026; Zhu, 2016), conventional bad-state measures need not identify the

periods in which marginal utility is highest.

We therefore consider a second set of candidate high-marginal-utility states based on uncertainty about the direction of government action. When the government's response is itself uncertain, households and investors cannot rely as confidently on the implicit protection that the government ordinarily provides. Such periods may therefore be especially likely to coincide with high marginal utility in China's discretion-based economic system. We refer to this uncertainty as "policy-regime risk."

We construct two policy-regime measures. The first uses three IMF-motivated episodes involving changes in the direction, credibility, or continuation of government policy: the mid-2015 equity correction and exchange-rate-policy change; the 2018 transition between deleveraging and renewed growth support; and the 2021 property-financing and regulatory tightening. The second uses four China-specific policy-direction episodes: the 2014 escalation of the anti-corruption campaign; the 2018 trade-war escalation with the United States; the 2020 introduction of property-financing rules; and the late-2020 shift in platform and technology regulation. Appendix G explains the selection and dating of these episodes.

We also employ two sets of conventional high-marginal-utility proxies. The first captures stock-market returns, market volatility, and economic policy uncertainty. The second captures consumer confidence, manufacturing demand, and household assessments of income and employment conditions. Appendix G provides the variable definitions.

Table 6 reports separate regressions of the small supported-minus-non-supported return spread on each candidate state measure, controlling for the CH4 factors. Panels A and B report the conventional financial-market and demand-state results. The coefficients are statistically insignificant throughout. The down-market, consumer-confidence, manufacturing-demand, and household-condition indicators have positive coefficients. Although these es-

timates are imprecise, the signs are consistent with the possibility that weak market and economic conditions increase expectations of government support. The economic policy uncertainty coefficients are negative but statistically insignificant.

Panel C reports the policy-regime results. The coefficients on both episode measures are negative and statistically significant under the one-sided tests reported in the table. The result is strongest for the China-specific policy-direction episodes, for which the estimated state coefficient is -0.013 with a p-value of 1.2%. Figure 5 plots factor-adjusted portfolio returns around these episodes. The IMF-motivated episodes are associated with a sharp decline at the beginning of the episode window. The China-specific episodes are associated with lower returns over approximately five months surrounding the policy-direction events. These results suggest that the policy portfolio performs poorly during periods of substantial uncertainty about the direction of government action.

Taken together, the conventional measures provide little evidence that the policy portfolio performs poorly during standard periods of market or economic stress. In China, however, the relationship of such periods with household marginal utility is ambiguous. By contrast, the portfolio performs poorly during policy-regime episodes in which the future direction of the country is uncertain. If policy-regime risk is a good proxy for periods of high-marginal-utility in China, then these documented negative correlations are consistent with the existence of a policy risk premium.

6.2 Heterogeneous Exposure to Policy-Support Uncertainty

In this section, we explore heterogeneous exposure to policy-support uncertainty. Equation 5 shows that the premium depends on both the value of policy support relative to firm value, $B_i/P_{i,t}$, and the state dependence of changes in expected support. This suggests that firms

more exposed to the policy risk should earn higher premiums.

We investigate differences between state-owned and non-SOE firms, differences in the likelihood of local implementation, and differences in firm size. In each setting, we find evidence consistent with the policy-risk interpretation.

6.2.1 State Ownership

State ownership provides a measure of exposure to policy-support uncertainty. SOEs already possess institutional connections to government agencies, state-controlled banks, and other public-sector resources. A new central policy directive may therefore change the expected resources available to an SOE less than it changes the prospects of a non-SOE. That is, the support benefit B_i may be relatively larger for non-SOEs than for SOEs. If the abnormal return compensates investors for uncertainty about whether policy support will translate into firm-level benefits, the return should consequently be stronger among non-SOEs.

We examine this prediction by independently sorting small firms on policy support and SOE status. Panel A of Table 7 presents the results. The value-weighted CH4 alpha is 0.940% per month for supported non-SOEs and 0.065% for supported SOEs. The directly estimated non-SOE-minus-SOE difference is 0.876 percentage points ($p = 0.004$). This ownership split is consistent with greater exposure to policy-support uncertainty among non-SOEs.

6.2.2 Central–Local Five-Year Plan Alignment

Central policy priorities rely on local governments for implementation. Local officials control or influence many of the land, credit, approvals, subsidies, and state-linked resources required to translate a central directive into firm-level benefits. We compare supported firms based on whether the province in which each firm is registered designates the firm's industry as a

priority in the applicable provincial Five-Year Plan.

The results are presented in Panel B of Table 7. The value-weighted monthly CH4 alpha is 0.829% for the local-FYP portfolio and 0.088% for the no-local-FYP portfolio. The directly estimated difference is 0.740 percentage points ($p = 0.014$).

This pattern is consistent with implementation requiring alignment across levels of government. When central and local priorities conflict, local implementation may be less likely. If local officials do not intend to implement a central policy, the expected benefit B_i, or the probability that this benefit will be realized, will be lower.[20]

6.2.3 Firm Size

Our baseline results (Table 5), which demonstrate that abnormal returns are concentrated among small firms, are also consistent with greater exposure to policy risk. Holding the benefit B_i fixed, the value of the benefit relative to firm value is larger when $P_{i,t}$ is smaller (see $B_i/P_{i,t}$).

Furthermore, Chinese industrial policies are often meant to spur competition in a given industry. Large incumbent firms, which may have previously enjoyed rents or support from local and central government officials, may not benefit as much if new policies increase competition. That is, the benefits may be both relatively and absolutely more important for small firms, and the discontinuation of policy support may thus disproportionately harm small firms.

Taken together, these heterogeneity tests are consistent with the policy-risk interpretation. Across the ownership, central–local alignment, and size tests, groups with greater exposure to policy-support uncertainty earn higher abnormal returns.

[20] An absence from a local plan need not represent political resistance: local officials may agree that the centrally favored industry is not an appropriate development priority for that particular region.

7 Alternative Interpretations and Robustness

Policy assignment is endogenous, but endogenous assignment is not by itself sufficient to generate predictable post-announcement returns. If the government's selection of an industry reveals information about future fundamentals, rational investors should incorporate that information when the policy becomes public. Government selection can produce persistent abnormal returns only if investors fail to incorporate the signal fully, market frictions delay price adjustment, or selection proxies for an omitted priced risk.

In this section, we therefore examine whether systematic mispricing, concentrated industry performance, or omitted risk factors can explain the abnormal returns. We also test whether the results are sensitive to how policy support is measured.

No individual test can establish that the benchmark asset-pricing models are complete or eliminate every form of persistent belief error. Our objective is instead to determine whether competing explanations can jointly accommodate the broad set of findings documented in the preceding sections.

7.1 Systematic Mispricing

Recall from equation 8 that expected factor-adjusted returns can be written as

$$E_t\left[\widetilde{R}_{i,t+1}\right] = \mu_{i,t}^{P} + \frac{B_i}{P_{i,t}} E_t\left[\Delta p_{i,t+1}\right] + E_t\left[\varepsilon_{i,t+1}\right]$$

The realized idiosyncratic term should average out in a sufficiently diversified portfolio, an assumption we examine in Section 7.2. If investor beliefs contain systematic errors, or if market frictions prevent prices from incorporating available information, realized belief

revisions may have a nonzero average. Subsequent abnormal returns could then reflect the correction of mispricing as well as compensation for policy-support risk.

The tests in this section example whether the return pattern exhibits the attention, exuberance, or arbitrage-friction characteristics predicted by this alternative explanation.

7.1.1 Investor Attention

One potential explanation for gradual price increases is that investors are inattentive and policy information diffuses slowly into prices over time (Hong and Stein, 1999). If this is the case, then the abnormal returns we observe may not reflect compensation for risk, but rather a delay in how quickly markets respond to public news.

This is a particularly important consideration because policies are endogenous. If the government "picks winners" and investors do not fully account for this, then the large abnormal returns we document may reflect endogenous policy with inattentive investors.

Investor inattention is unlikely to explain our results. The industrial policies we study are publicly announced and typically follow lengthy public-private deliberations that receive substantial media attention. To test this formally, we first examine whether investors, including retail investors, are likely to be aware of these policies, and then we examine whether investors and the public are paying attention to these industries. We also show that our main results are not driven by low-attention policies or industries.

In Appendix H, we analyze transcripts from *Xinwen Lianbo*, CCTV's flagship daily news program and an important dissemination channel for Chinese government announcements and Communist Party positions (Shao et al., 2015). We find that 58.5% of the supportive policies in the pricing sample are discussed within three days of publication, and we classify these as having "rapid coverage."

Table H1 compares portfolios of small supported stocks in industries with and without rapid coverage. The rapid-coverage portfolio includes industries with at least one rapidly covered supportive document during the preceding 12 months. Monthly CH4 alphas for this portfolio are 0.63% under both equal and value weighting, significant at the 1% level. The no-rapid-coverage alphas are small, positive, and statistically insignificant. The unmentioned policies do not earn statistically significant alpha, which is inconsistent with the investor-inattention hypothesis.

We also examine two measures of industry-level attention: (1) public attention, captured by the relative volume of industry-specific news coverage,[21] and (2) investor attention, proxied by industry-level trading volume relative to total market trading. Both are computed daily and regressed on event-time indicators to detect abnormal attention levels around policy issuance dates.

Figure 6 shows that attention is high in the year leading up to and the six months following the issuance of supportive policy documents. News coverage (Panel A) is elevated well in advance of policy issuance, consistent with anticipation of support, and trading activity (Panel B) is elevated both before and after policy issuance.

Under the inattention hypothesis, information about policy support moves slowly and we would expect abnormal returns to be concentrated among firms receiving low attention, since informed trading would be delayed. Table J4 of Appendix J shows the opposite pattern. When using public attention (Panel A), differences are statistically insignificant. When using investor attention (Panel B), the difference in returns between the low- and high-attention groups is consistently negative and mostly significant. Abnormal returns are therefore not

[21] News data are obtained from the Financial News Database in CSMAR and cover financial news published by Chinese mainstream media outlets. We have data for 329 industry-policy combinations between March 28, 2008, and June 30, 2015.

concentrated in low-attention industries, as a simple slow-information-diffusion explanation would predict.

7.1.2 Extrapolative Beliefs

High investor attention raises the possibility that our results are driven by another common mispricing mechanism: extrapolative beliefs (Barberis et al., 1998). Investors with such beliefs extrapolate past price movements and purchase these stocks under the belief that prices will continue to rise. This buying pressure raises prices and may cause the excess returns we observe. Retail investors, who represent a large fraction of trading activity in the domestic Chinese stock market, are known to exhibit extrapolative expectations (Liao, Peng, and Zhu, 2022).

If policy implementation is genuinely valuable, then stock prices should rise following the issuance of a supportive policy document. However, investors with extrapolative beliefs may expect these price increases to continue and thus purchase these stocks, pushing prices excessively high. This buying pressure generates returns in the short run that are not fundamentally justified.

CSMAR provides daily margin-trading data at the stock level, which we use as a proxy for leveraged speculative demand. If exuberant investors push prices above fundamentals, margin-financed trading should increase around supportive-policy announcements. Appendix H.2 analyzes margin financing around these events and finds no unusual increase.

Nevertheless, the abnormal returns we document may reflect investor exuberance. In this case, investors push prices too high, but prices should eventually revert and abnormal returns should turn negative in later periods. Following the approach of Jegadeesh and Titman (1993), we test this hypothesis by examining the long-term performance of supported

stocks. Using information at month t to identify supported stocks, we analyze future returns between month $t+k-1$ and $t+k$, for k ranging from 1 to 36. Our main analysis in Section 5.2 focuses on $k = 1$, and a reversal would manifest as negative abnormal returns for some $k > 1$.

Figure 7 plots the alphas over time. Across all three models, the estimated alphas decline gradually but remain positive throughout the post-policy window. This pattern is inconsistent with return reversals and suggests that investors are not systematically over-extrapolating from policy announcements.

7.1.3 Limits to Arbitrage

Even if investors hold rational beliefs, mispricing may still arise if there are substantial limits to arbitrage (Shleifer and Vishny, 1997). We examine whether several prominent limits to arbitrage could sustain the return pattern.

First, our investment strategy does not rely on short-selling. Short-selling is costly in China, and the government has at times imposed implicit or explicit short-sale restrictions during periods of market stress. However, our strategy can be implemented through simple long positions in supported stocks using a standard buy-and-hold approach. The strategy therefore avoids the direct costs and regulatory constraints associated with short-selling.

Second, while our results are concentrated among small firms, we find little evidence that these abnormal returns are driven by illiquidity. We reduce the influence of illiquid stocks in three ways: (1) the baseline sample requires a minimum trading history and at least 15 trading days in the preceding month; (2) value weighting reduces the influence of the smallest stocks; and (3) the CH4 factors are constructed after excluding the smallest 30% of firms by market capitalization.

Third, the investment strategy does not require proprietary information. The underlying policy documents and industry classifications are publicly available, so the strategy does not depend on private data, although processing and mapping the documents requires effort.

Taken together, the evidence is difficult to reconcile with simple versions of investor inattention, extrapolative beliefs, or prominent arbitrage constraints. These tests cannot rule out every form of persistent mispricing, but the observed return pattern raises the evidentiary burden for a systematic-mispricing explanation.

7.2 Industry Concentration and Economic Magnitude

The supported portfolio contains firms from many industries, but diversification across stocks may not eliminate the influence of concentrated industry performance. A small number of industries may receive repeated support and realize unusually high returns over our sample. In that case, the realized sample average of $\varepsilon_{i,t+1}$ in equation 8 may not converge to zero. That is, our headline alpha result may partly reflect favorable ex post industry outcomes rather than a return pattern that extends broadly across supported sectors. We evaluate this concern by excluding economically prominent groups and by measuring the sensitivity of abnormal returns to the best-performing industries.

Technology is the most salient potential source of concentration. China's 13th Five-Year Plan and related initiatives placed substantial emphasis on innovation and high-technology development during the middle of our sample, a period that also contained strong performance by Chinese technology firms. We therefore remove technology-sector firms and re-estimate the supported portfolio. Panel A of Table 8 shows that the value-weighted portfolio of small supported firms continues to earn a CH4 alpha of approximately 0.5% per month, or 6.0% per year. The headline result is therefore not confined to the technology

sector.

We next impose a more demanding ex post concentration test. We rank industries by realized performance over the full sample, remove the top decile, and reconstruct the supported portfolio. The results are reported in Panel A of Table 8. This brute-force test reduces the annualized CH4 alpha of small supported firms from 7.2% to approximately 3.6%. The estimate remains positive and significant at the 5% level. This is our most stringent concentration test and produces a conservative annualized estimate of approximately 3.6%.

These exercises measure the extent to which unusually successful industries contribute to the headline results. The remaining alpha suggests that unusually successful industries do not account for the entire result.

7.3 Economic Policy Uncertainty and R&D Risk

The portfolio alpha may represent compensation for a source of expected returns that is correlated with policy support but omitted from our benchmark models. For such an explanation to account for the results, the omitted mechanism must command a premium, covary with the supported portfolio, remain outside the CAPM, Carhart, and CH4 factors, and explain the magnitude and structure of the alpha. The stability of the alpha across three established benchmarks reduces this concern but cannot establish that any factor model is complete. We focus on two economically plausible omitted risks: aggregate economic policy uncertainty and R&D investment.

The first is aggregate economic policy uncertainty (EPU). We use two China-specific EPU indices: the *South China Morning Post* measure of Baker et al. (2016) and the *Renmin Daily* and *Guangming Daily* measure of Davis et al. (2019). Following Brogaard and Detzel

(2015), we extract innovations in each index and construct factor-mimicking portfolios from Chinese equity returns. The full construction is presented in Appendix I. The first two rows of Panel B in Table 8 estimate CH4 alpha in a model that has been extended to include these EPU factors. The supported-portfolio alpha remains economically and statistically significant under both specifications and is not materially smaller than the corresponding benchmark estimate. Broad uncertainty about fiscal, regulatory, monetary, and trade policy therefore does not absorb the return associated with industry-specific support.[22]

The second candidate is R&D investment risk. Section 4.1 shows that stronger support is followed by greater R&D spending, so the supported portfolio may load on an innovation premium unrelated to policy-support uncertainty. Following Leung et al. (2020), we construct an R&D factor from value-weighted portfolios sorted on firm size and R&D intensity (see Appendix I). The final row of Panel B in Table 8 shows that adding this factor leaves the CH4 alpha on the small supported portfolio largely unchanged. The existing results therefore do not indicate that exposure to the constructed R&D factor accounts for the abnormal return.

The existing EPU and R&D tests show that these two plausible omitted factors do not absorb the estimated abnormal return. These tests do not rule out every possible omitted source of expected returns, but neither EPU nor the constructed R&D factor explains the abnormal return.

[22]The EPU measures may combine states with different implications for policy risk. EPU can be high during market instability, when government intervention is expected, and during periods when the direction of government policy is itself uncertain.

7.4 Policy Measurement and Specification Robustness

Our baseline portfolio depends on classifying State Council documents as supportive and mapping policy themes to listed-firm industries. Because policy language ranges from broad encouragement to concrete resource allocation, classification error or reliance on particular title words could affect the return estimate. We therefore examine whether the result survives alternative human- and model-based classifications, exclusions of influential keywords, and economically meaningful changes to the portfolio definition.

Our baseline classification scheme uses title keywords to identify supportive policies. In Appendix A.3, we describe two additional classification schemes. One is based on manually reading each document and assigning a support indicator. The other asks ChatGPT-3.5 to read the documents and determine which documents are supportive. As mentioned in Section 3.2, there is a high degree of agreement across these methods.

Panel C of Table 8 reports the CH4 alpha for small supported firms under manual classification, GPT-based classification, and the intersection of the baseline keyword-based measure with these additional measures. Under all three measures, the alpha remains positive and statistically significant, suggesting that our results are not sensitive to these alternative classification schemes.

We next examine whether our results are sensitive to influential words in our baseline title-keyword classification scheme. Panel A of Table 9 presents two analyses. The first excludes policies triggered only by the generic terms *reform*, *plan*, and *strengthen*; the second excludes policies triggered by the three largest sole-trigger words *strengthen*, *accelerate*, and *reform*. In both cases, the CH4 alpha remains positive and significant, suggesting that generic or influential keywords are not the sole driver behind our classification scheme.

Finally, we use Chinese-language DeepSeek models to extract policy direction, concrete measures, implementation details, resource channels, and supporting text (see Appendix A.4). The analysis covers all 404 documents and 1,074 policy–industry links. Repeated applications show high within-model agreement, which measures repeatability rather than classification accuracy. Panel B of Table 9 reports monthly CH4 alphas across the Pro- and Flash-model variants; all are significant at the 1% level.

As discussed in Section 3.2, our DeepSeek algorithm also classifies policies based on the intensity of support. The categories include (1) vague guidance, (2) executable nonresource support, and (3) resource and/or cash-flow support. Panel C of Table 9 shows that most documents classified as supportive by DeepSeek provide resource or cash-flow support. The resource-support portfolio produces a positive and statistically significant CH4 alpha, while the executable-nonresource and vague-guidance portfolios do not. The direct differences across intensity categories are imprecisely estimated, so the evidence is suggestive rather than a definitive monotonic test of policy intensity.

Taken together, these results suggest that our analysis is not sensitive to the baseline keyword-based classification of supportive policies used in our main analysis.

8 Conclusion

This paper studies how investors price Chinese industrial-policy risk. Using a new dataset of central government directives, we show that policy support is associated with meaningful changes in firm resources, operations, investment, and downside risk. These benefits, however, are uncertain and depend on the implementation and continuation of support. Small firms in supported industries earn annualized abnormal returns of approximately 7.2% rela-

tive to our preferred China-specific asset-pricing benchmark.

The state dependence and cross-sectional structure of these returns are consistent with compensation for policy-support risk. The policy portfolio performs poorly when the direction of government action is uncertain, and abnormal returns are higher among firms with greater exposure to policy-support uncertainty. The return pattern is also difficult to reconcile with simple explanations based on investor inattention, extrapolative beliefs, limits to arbitrage, or concentrated industry performance.

More broadly, our analysis highlights how political institutions shape economic risk. Under China's discretion-based system, the central government "leads everything" through its policies. Because the government has the power to dynamically change these policies, uncertainty about the direction of government action becomes a persistent feature of the economic environment. More fundamentally, the states in which marginal utility is highest may include those in which households and investors are especially uncertain about the intentions of central decision makers and, therefore, about the future direction of the country. The broader lesson is that institutions shape both the risks economic actors bear and also the states of the world in which those risks matter most.

References

Aghion, P., C. Jing, M. Dewatripont, L. Du, A. Harrison, and P. Legros (2015). Industrial policy and competition. *American Economic Journal: Macroeconomics 7*(4), 1–32.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of Financial Economics 49*(3), 307–343.

Barwick, P. J., M. Kalouptsidi, and N. B. Zahur (2019). China's industrial policy: An empirical evaluation. Technical report, National Bureau of Economic Research.

Brogaard, J. and A. Detzel (2015). The asset-pricing implications of government economic policy uncertainty. *Management Science 61*(1), 3–18.

Brunnermeier, M. K., M. Sockin, and W. Xiong (2017). China's gradualistic economic approach and financial markets. *American Economic Review 107*(5), 608–613.

Brunnermeier, M. K., M. Sockin, and W. Xiong (2022). China's model of managing the financial system. *The Review of Economic Studies 89*(6), 3115–3153.

Cai, H., H. Fang, and L. C. Xu (2011). Eat, drink, firms, government: An investigation of corruption from the entertainment and travel costs of Chinese firms. *The Journal of Law & Economics 54*(1), 55–78.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance 52*(1), 57–82.

Cheung, S. N. (2014). The economic system of China. *Man and the Economy 1*(1), 1–49.

Davis, S. J., D. Liu, and X. S. Sheng (2019). Economic policy uncertainty in China since 1949: The view from mainland newspapers. *Working Paper*.

Gebhardt, W. R., C. M. C. Lee, and B. Swaminathan (2001). Toward an implied cost of capital. *Journal of Accounting Research 39*(1), 135–176.

Hong, H. and J. C. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance 54*(6), 2143–2184.

Hou, K., M. A. van Dijk, and Y. Zhang (2012). The implied cost of capital: A new approach. *Journal of Accounting and Economics 53*(3), 504–526.

Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance 48*(1), 65–91.

Kalouptsidi, M. (2018). Detection and impact of industrial subsidies: The case of Chinese shipbuilding. *The Review of Economic Studies 85*(2), 1111–1158.

Kelly, B., L. Pástor, and P. Veronesi (2016). The price of political uncertainty: Theory and evidence from the option market. *The Journal of Finance 71*(5), 2417–2480.

Kydland, F. E. and E. C. Prescott (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of political economy 85*(3), 473–491.

Leung, W. S., K. P. Evans, and K. Mazouz (2020). The R&D anomaly: Risk or mispricing? *Journal of Banking & Finance 115*, 105815.

Li, H., L. Meng, Q. Wang, and L.-A. Zhou (2008). Political connections, financing and firm performance: Evidence from Chinese private firms. *Journal of Development Economics 87*(2), 283–299.

Liao, J., C. Peng, and N. Zhu (2022). Extrapolative bubbles and trading volume. *The Review of Financial Studies 35*(4), 1682–1722.

Liu, A. Y. (2023). Beijing's banking balloon: China's core economic challenge in the new era. *The Washington Quarterly 46*(2), 69–86.

Liu, E. (2019). Industrial policies in production networks. *The Quarterly Journal of Economics 134*(4), 1883–1948.

Liu, J., R. F. Stambaugh, and Y. Yuan (2019). Size and value in China. *Journal of Financial Economics 134*(1), 48–69.

Naughton, B. (2021). *The Rise of China's Industrial Policy, 1978 to 2020*. Lynne Rienner.

Pástor, L. and P. Veronesi (2012). Uncertainty about government policy and stock prices. *The Journal of Finance 67*(4), 1219–1264.

Pástor, L. and P. Veronesi (2013). Political uncertainty and risk premia. *Journal of Financial Economics 110*(3), 520–545.

Shao, Z., X. Zhang, and T. Meng (2015). An analysis of the propaganda model of *Xinwen Lianbo* from a political communication perspective. *Tsinghua University Journal (Philosophy and Social Sciences) 30*(3), 30–42. In Chinese.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Sinclair, A. J. (2026). Implicit government guarantees. *Working Paper*.

Sun, G. (2020). Banking institutions and banking regulations. *The Handbook of China's Financial System*, 9–37.

Xu, C. (2011). The fundamental institutions of China's reforms and development. *Journal of Economic Literature 49*(4), 1076–1151.

Xu, C. (2025). *Institutional Genes: Origins of China's Institutions and Totalitarianism*. Cambridge University Press.

Zhu, N. (2016). *China's Guaranteed Bubble*. McGraw-Hill Education.

Figure 1: **Number of Supportive Industrial Policy Documents Published Each Year**
This figure shows the number of supportive industrial policy documents issued annually by the State Council from April 2008 to May 2022. In total, 282 documents are classified as "supportive," meaning that these documents explicitly direct local officials to promote targeted industries.

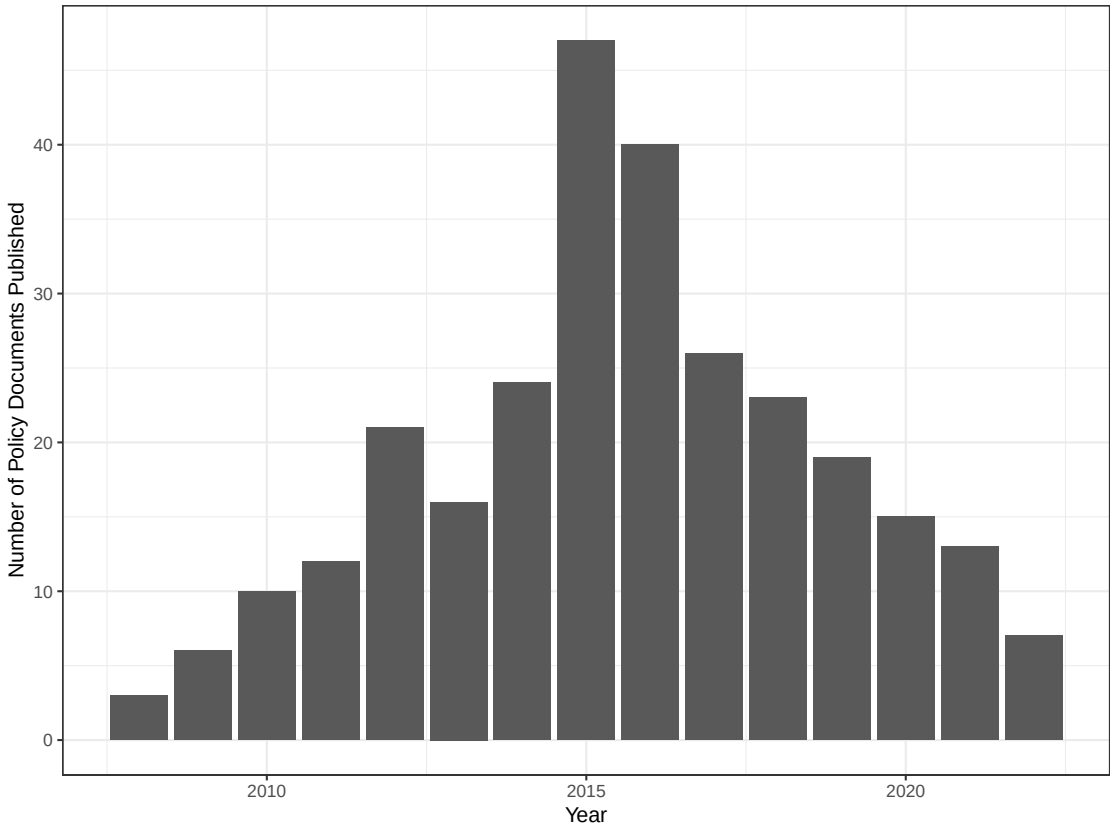
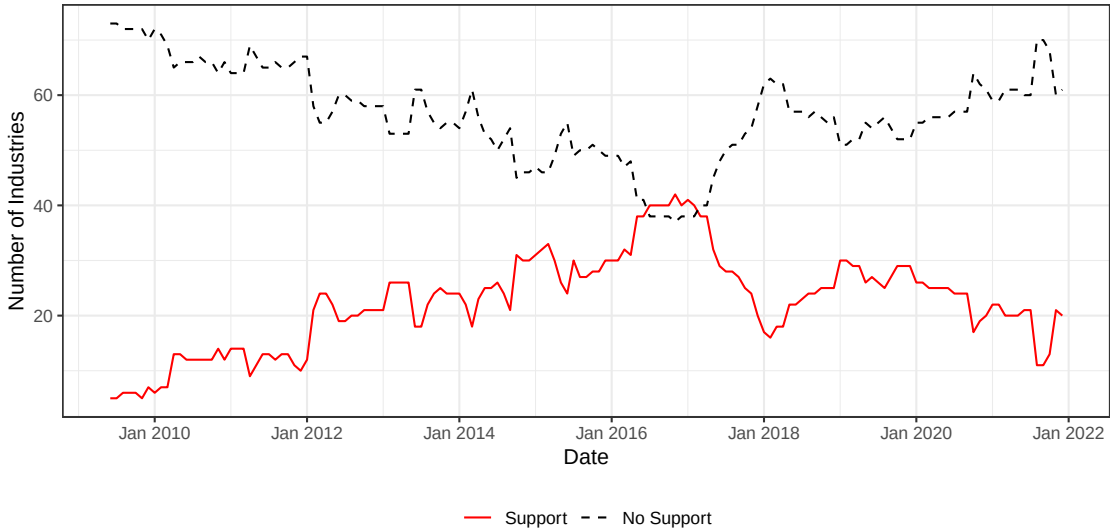


Figure 2: **Number of Centrally Supported Industries and Stocks**
This figure presents the number of industries (Panel A) and publicly listed stocks (Panel B) that received central government policy support over the preceding 12 months. An industry is classified as supported if, within the past year, the State Council has issued at least one policy document explicitly promoting its development. Stock data are obtained from the CSMAR database and include all A-share listings on the Shanghai, Shenzhen, and GEM boards. The analysis period runs from May 2009 to December 2021.

Panel A: Number of Centrally Supported Industries



Panel B: Number of Centrally Supported Stocks

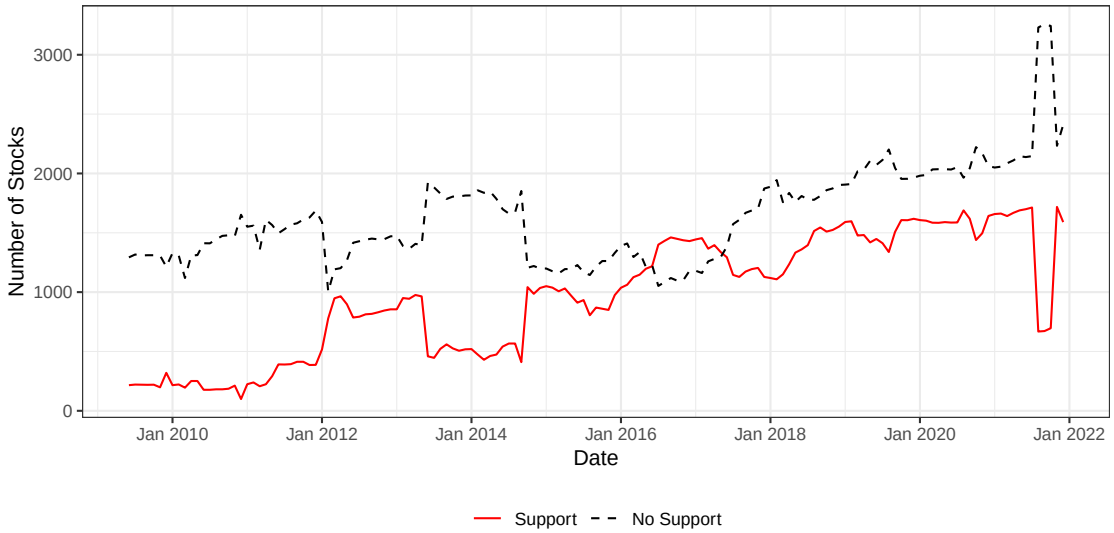
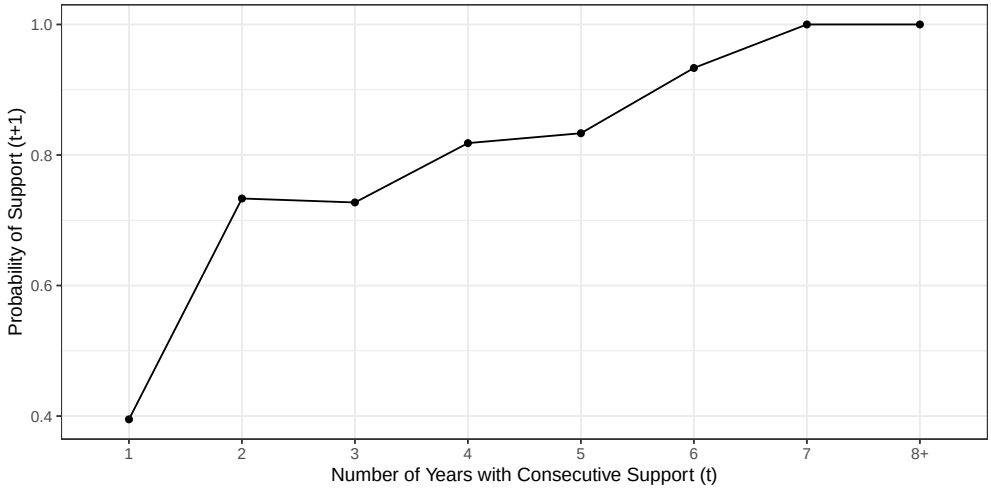


Figure 3: **Duration and Dynamics of Policy Support**
This figure illustrates the duration dynamics of industrial policy support across industries. Panel A plots the probability that support continues into year $t + 1$, conditional on having received consecutive support through year t. The probability of continuation rises as support accumulates, highlighting path dependence. Panel B shows the distribution of industries by their maximum streak of consecutive support years. The pattern is sharply bimodal: most industries receive either brief support (one year) or prolonged support (eight years or more). The sample includes policy documents issued between April 24, 2008, and May 27, 2022.

Panel A: Conditional Probability of Continued Support



Panel B: Distribution of Maximum Years of Consecutive Support

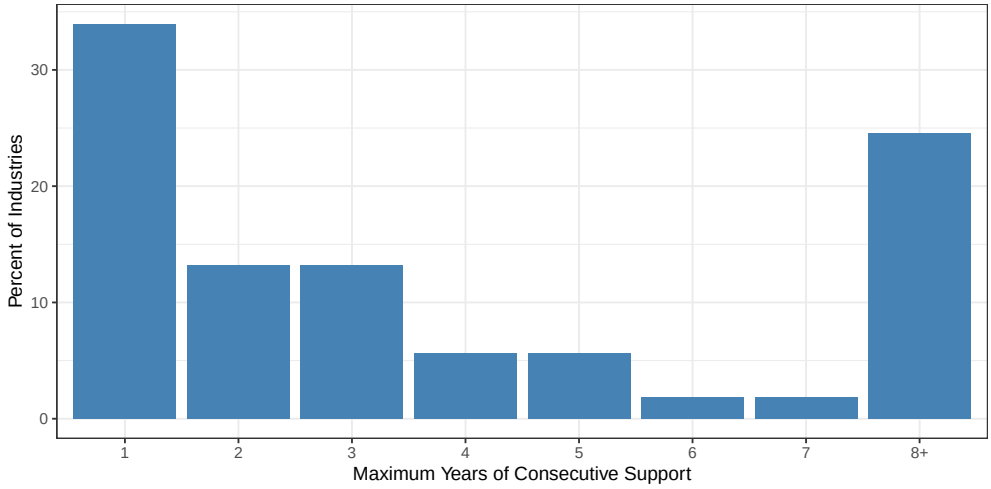
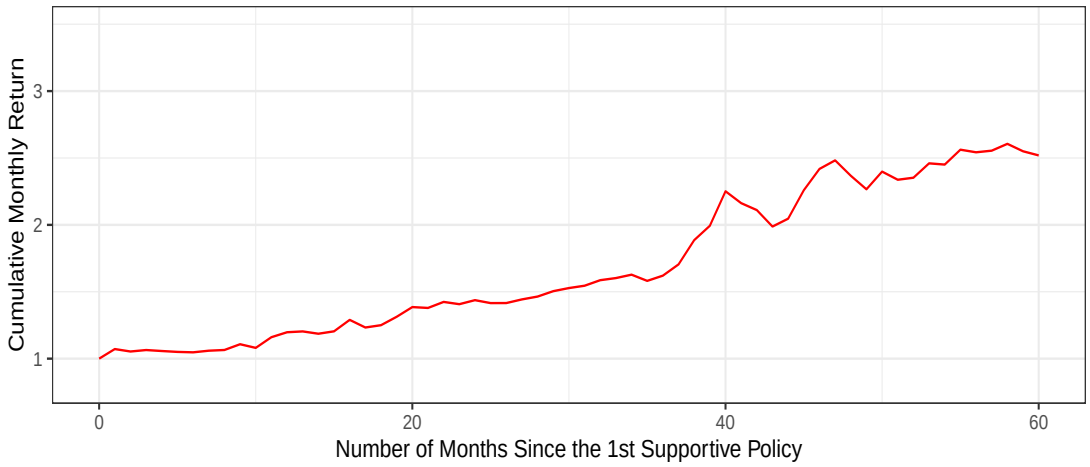


Figure 4: **Stock Returns Following Continued and Discontinued Support**
Panels A–C present equal-weighted firm-level stock returns in event time following the first supportive central government policy document for an industry. Panel A plots gross-return indexes, normalized to one at $t = 0$, through month 60. Panel B classifies initial episodes ex post according to whether the industry receives only one year of support (dashed black line) or at least two consecutive years of support (solid red line) and plots each group's gross-return index with the same normalization. Panel C reports average monthly returns and 95% confidence intervals in decimals for the two groups during each of the five years following initial support. First-year returns are positive and statistically indistinguishable, while the paths subsequently diverge. Because renewal is not known at the initial document date, Panels B and C characterize the ex post resolution of support rather than returns to an implementable initial-support strategy. Panel D reports returns following the last supportive policy document and is described on the continued page.

Panel A: Average Returns Following Initial Policy Support



Panel B: Returns Diverge Based on Duration of Support

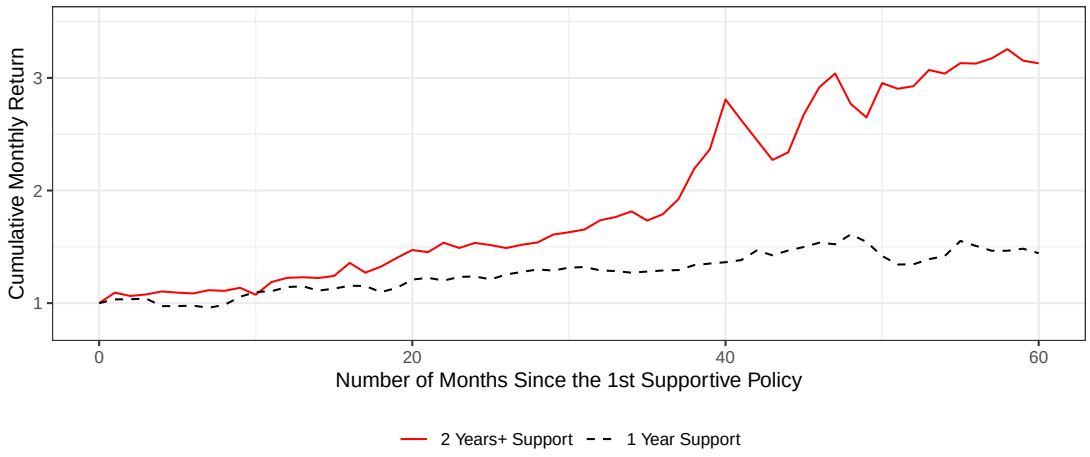
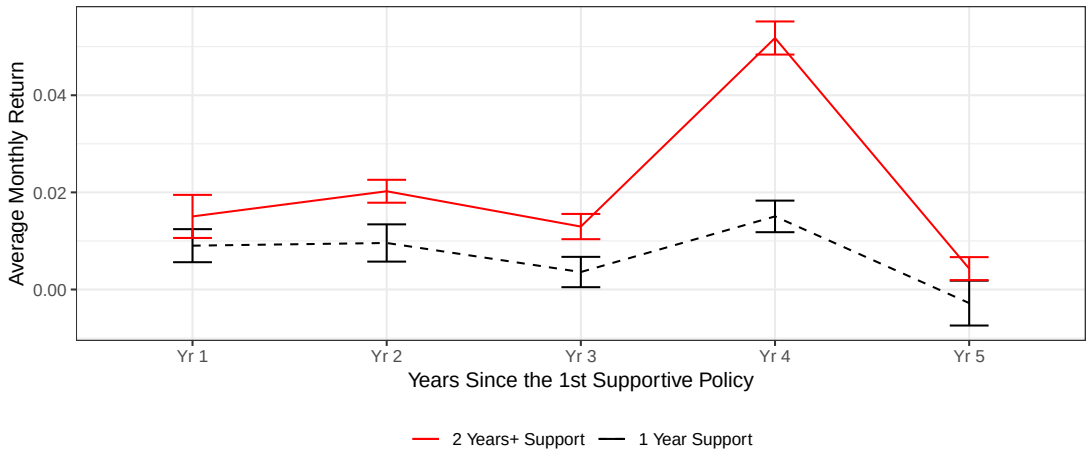


Figure 4: **Stock Returns Following Continued and Discontinued Support (continued)** Panel D illustrates that average returns drop when policy support disappears. It plots the average monthly returns and 95% confidence intervals for firms in the three years following the issuance of the last supportive central government policy document. An industry is defined as receiving its last supportive policy document if no subsequent document is issued within the next 36 months or if the policy is the final one for that industry during our sample period. Returns are presented in decimal form (e.g., 0.01 corresponds to 1.0% per month). The figure shows that returns are elevated in the first year, when investors believe that continued support is possible, but decline in subsequent years as it becomes clear that no further support is forthcoming. For reference, the average return on the market during this period is 0.73% per month, suggesting that returns are below the market average two and three years following the last supportive document.

Panel C: Average Monthly Returns Based on Duration of Support



Panel D: Average Returns Following the Last Supportive Policy Document

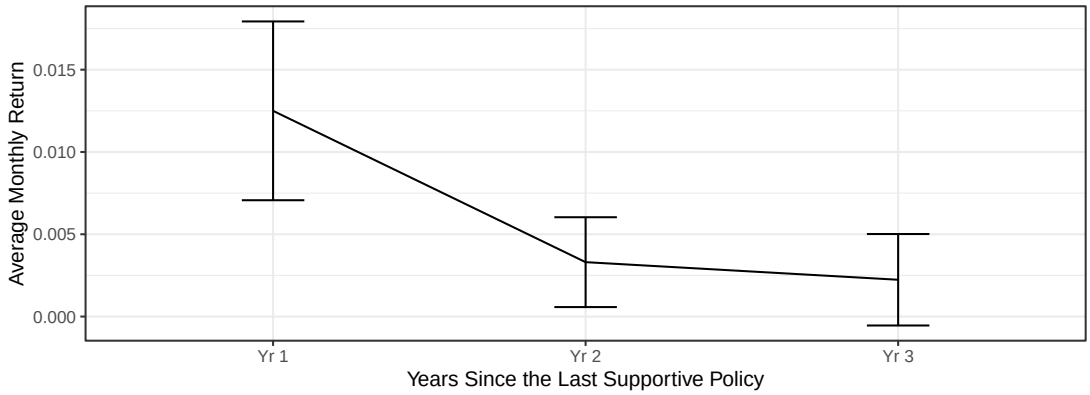


Figure 5: **Policy-Portfolio Returns Around China-Specific and IMF-Motivated Policy-Regime Episodes**

This figure shows event-time returns to supported-minus-non-supported small-firm portfolios around two sets of policy-regime episodes. Each panel plots average residuals from full-sample CH4 regressions for the equal- and value-weighted portfolios from event month −3 through +5. Residuals are monthly decimal returns. The shaded interval covers event months 0 through 2. Panel A averages the residuals around four China-specific policy-direction episodes. Panel B averages them around three externally dated IMF-motivated policy-regime episodes.

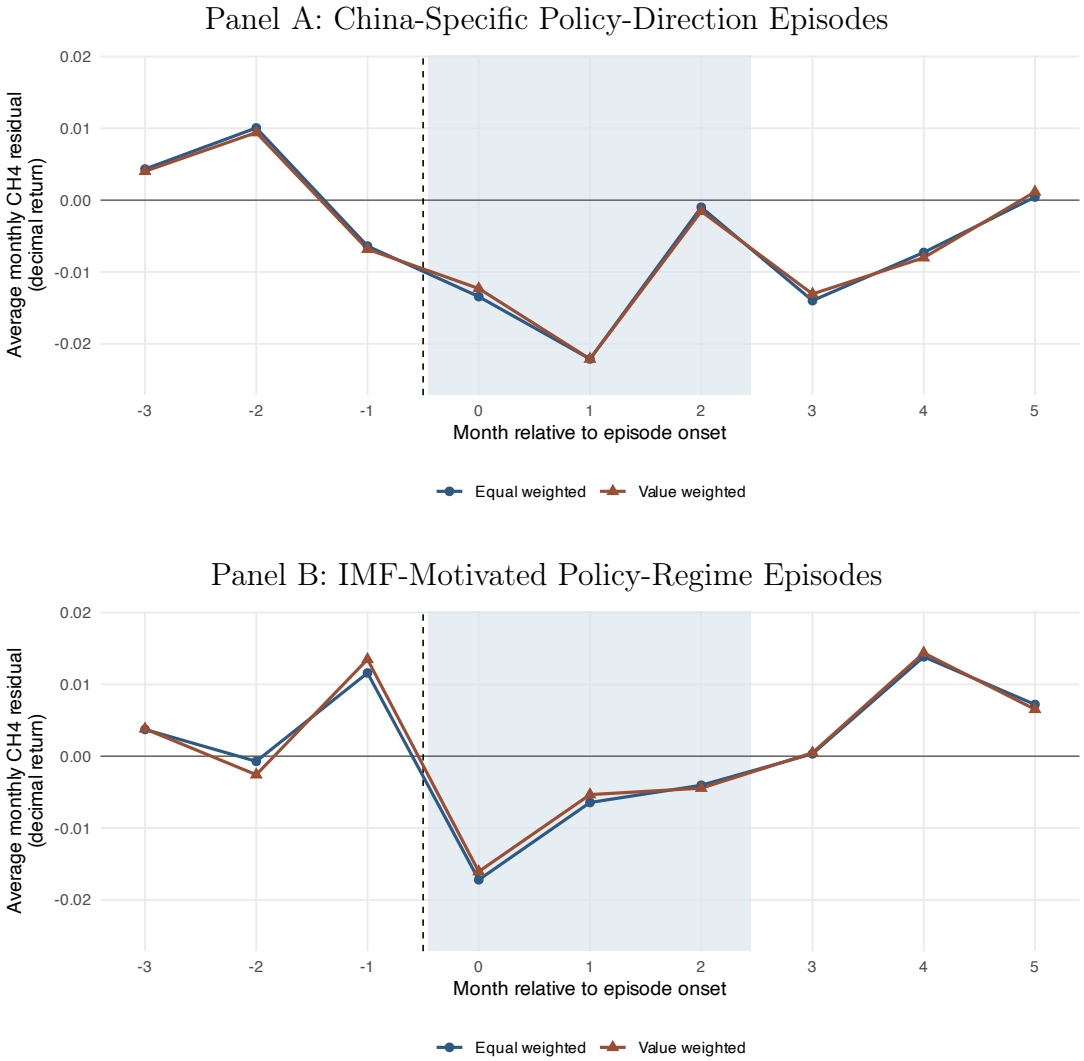
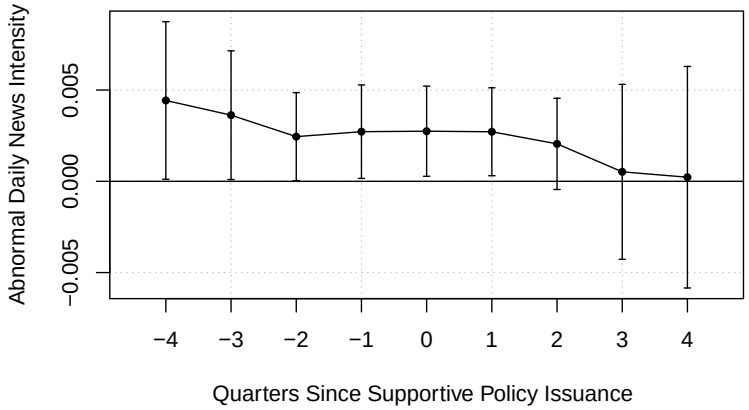


Figure 6: **Investor Attention Around Policy Issuance Date**
This figure tests whether investors are inattentive around the issuance of supportive industrial policies. Each panel plots abnormal investor attention in the four quarters before and after a policy announcement, using event-time coefficients and 95% confidence intervals from the following panel regression: $Attention_{i,t,q} = \sum_{k=-4}^{+4} \beta_k D^k_{i,t,q} + a_i + a_t + \epsilon_{i,t,q}$, where $D^{-4}_{i,t,q}$ to $D^{+4}_{i,t,q}$ are nine indicator variables for the quarters *relative* to the policy announcement, and a_i and a_t are industry and day fixed effects. In Panel A, attention is measured by news intensity, defined as the share of daily news articles related to industry i. In Panel B, attention is measured by trading volume share, defined as the fraction of daily market trading attributable to industry i. Standard errors are clustered at the industry-day level. The trading volume analysis covers the full sample, whereas the news coverage analysis, because of data availability, includes only industry-policy pairs between March 28, 2008, and June 30, 2015. In both cases, we find no evidence of investor inattention: attention increases prior to and following policy announcements, consistent with anticipatory and responsive behavior.

Panel A: Abnormal News Coverage Around Policy Announcements



Panel B: Abnormal Trading Volume Around Policy Announcements

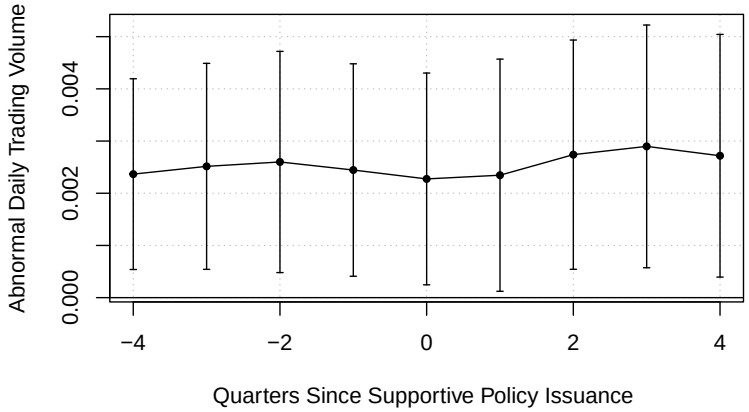


Figure 7: **Reversal Test: Long-Term Performance of Supported Firms**
This figure tests whether the abnormal returns to supported firms reverse in the long run. Specifically, we report monthly alphas and 95% confidence intervals in decimals for value-weighted portfolios of small supported stocks, constructed using information available at time t. For each k from 1 to 36, we compute the portfolio return from $t + k - 1$ to $t + k$, covering up to three years of forward performance. Panel A reports alphas estimated using the CAPM model; Panel B uses the Carhart four-factor model; and Panel C uses the CH4 model. All estimates are based on monthly data between May 2009 and December 2021. Across all specifications, abnormal returns remain positive over time, providing no evidence of return reversal.

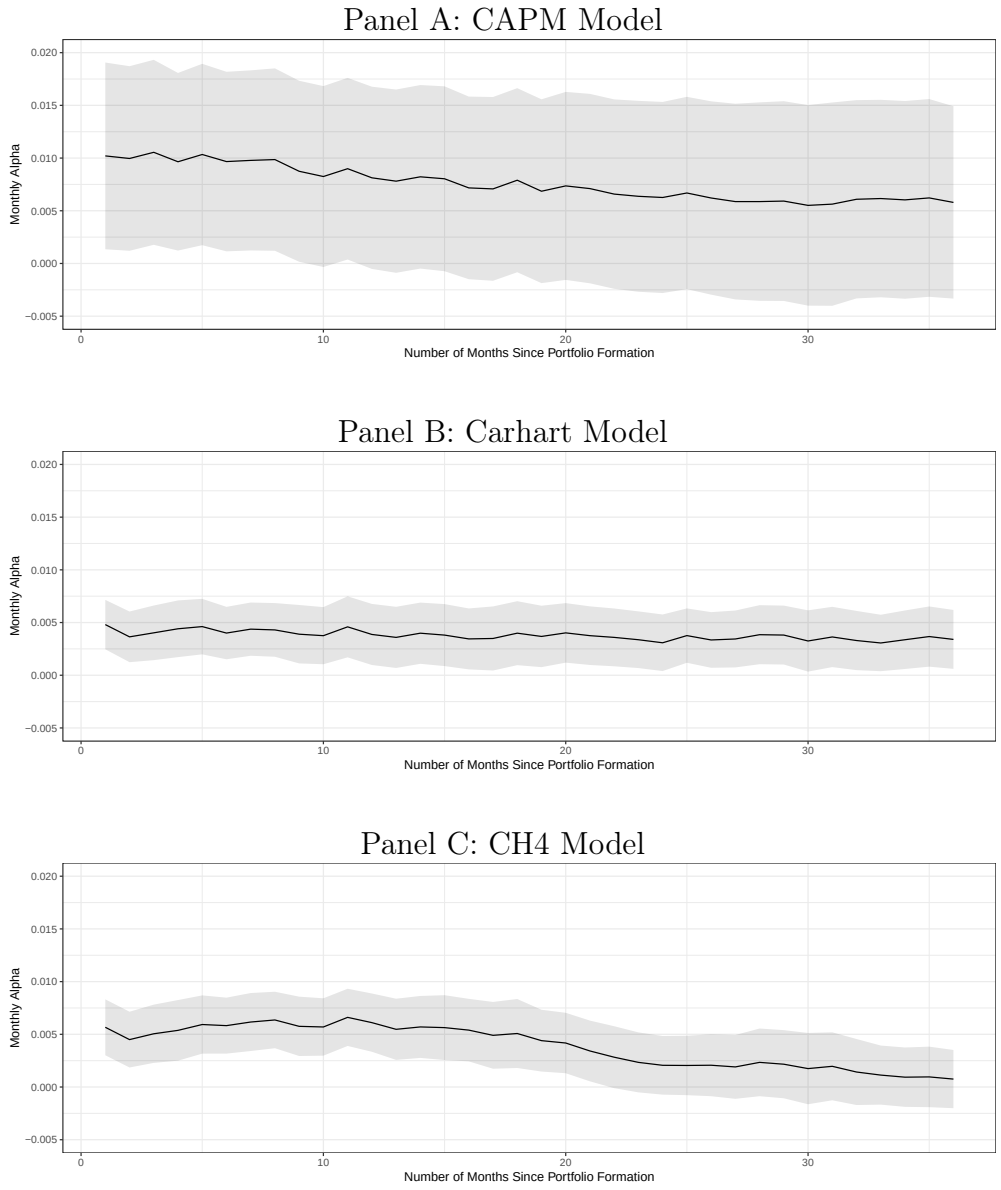


Figure 8: **Cumulative Excess Return of Supported Firms**
Consistent with the risk premium hypothesis, this figure demonstrates that the excess returns of supported firms are consistently earned throughout the sample period, except when the tech crackdown abruptly reversed policy in November 2020. Specifically, this figure plots the cumulative monthly CH4 alpha, expressed in decimals, earned by a value-weighted portfolio of small supported firms.

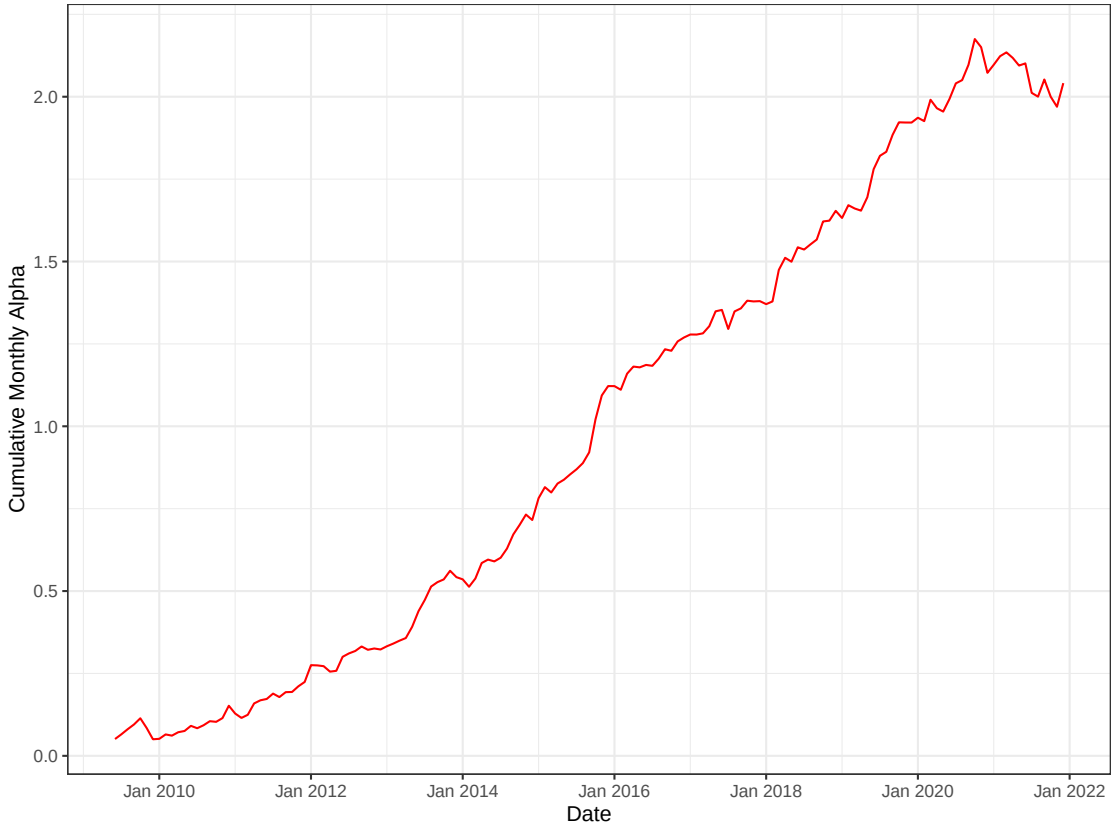


Table 1: **List of Policy Document Themes and Industry Mapping**

This table presents the 95 themes that the State Council database uses to classify policy documents. Panel A reports the themes, and asterisks (*) denote themes that can be mapped to one or more industries. Panel B reports the mapping between the themes and the 2012 CSRC Industry Classification of Listed Companies. The themes "Technology" and "Intellectual property" are mapped based on the definition of high-tech industries by the National Bureau of Statistics of China. The classification of high-tech industries is a derivative classification of the National Economic Industry Classification and refers to the OECD (Organisation for Economic Co-operation and Development) classification standards for high-tech industries.

Panel A: Themes

Administrative division and place name	Financial	∗Petroleum and natural gas
Administrative inspection	Food and drug supervision	∗Physical education
∗Aerospace	∗Foreign economic and trade cooperation	Population and family planning
∗Agriculture, animal husbandry, fisheries	∗Forestry	∗Postal
Arms resettlement	General office of the State Council	Poverty alleviation
Audit	Government publicity	Preferential resettlement
∗Bank	∗Health	Price
∗Broadcasting, movie, TV	∗Highway	Production safety supervision
∗Chemical	Hong Kong and Macau affairs	Public security
City planning	Human resource	Quality supervision
∗Civil aviation	Industry and commerce	∗Railway
∗Coal	∗Information industry (including telecommunications)	Religious affairs
Community management	Inspection	Secretarial work
Comprehensive use of energy conservation and resources	Inspection and quarantine	∗Securities
Confidentiality	∗Insurance	Social credit system construction
Cultural relic	∗Intellectual property	Social security
∗Culture	Judicial	Social welfare
Currency (including foreign exchange)	Labor	Standard
∗Customs	Land	State-owned asset supervision
Defense construction	∗Light textile	State Council
Defense industry	Literature and history	State Council's composition department
Diplomatic and foreign affairs	Macroeconomic	Statistics
Directly affiliated organs of the State Council	Major construction project	Taiwan affairs
Disaster reduction	∗Mechanical manufacturing and heavy industry	Tax
∗Domestic trade (including supply and marketing)	∗Medical management	∗Technology
E-government	Meteorological, hydrology, surveying and mapping, earthquake	Travel
Economic reform	∗Mineral	∗Urban and rural construction (including housing)
∗Education	Ministries and commissions of the State Council	∗Water conservancy
∗Electricity	National security	∗Water resources
Emergency management	∗News publishing	∗Water transport
Environmental monitoring, protection and governance	Ocean	Women and children
Ethnic affairs	Other	

Panel B: Theme-to-Industry Mapping	
Theme	CSRC Industries Mapped
Aerospace	56
Agriculture, animal husbandry, fisheries	4
Bank	66
Broadcasting, movie, TV	86
Chemical	7,11,25,45
Civil aviation	56
Coal	6,11
Culture	87
Customs	51
Domestic trade (including supply and marketing)	1,51,52
Education	82
Electricity	44
Foreign economic and trade cooperation	51
Forestry	2
Health	83
Highway	54
Information industry (including telecommunications)	63,64,65
Insurance	68
Intellectual property	27,35,39,40,63,64,65,73,74,75,77
Light textile	17,18
Mechanical manufacturing and heavy industry	32
Medical management	27
Mineral	9,32
News publishing	85
Petroleum and natural gas	7
Physical education	88
Postal	60
Railway	53
Securities	67
Technology	27,35,39,40,63,64,65,73,74,75,77
Urban and rural construction (including housing)	70
Water conservancy	76
Water resources	46
Water transport	55

Table 2: **Summary Statistics**
This table presents the summary statistics for our sample of 4,288 listed firms. These values are annual, are winsorized at the 1% level, and cover 2008–2021. SupportDummy is the main variable used to sort supported stocks and indicates whether a given firm has received a supportive policy document over the past 12 months.

Variable Name	Mean	Std.	10th percentile	Median	90th percentile
Firm Characteristics					
TotalAssets(Bil.)	9.236	23.529	0.643	2.579	18.142
Age	113.105	86.122	15.000	94.000	247.000
Leverage	0.415	0.212	0.144	0.403	0.689
Tangibility	0.230	0.170	0.034	0.197	0.481
OperatingRevenue(Bil.)	5.425	13.897	0.228	1.312	11.061
PoliticalConnection	0.402	0.472	0.000	0.000	1.099
EffectiveTaxRate	0.172	0.146	0.000	0.149	0.328
SupportStr	0.045	0.070	0.000	0.000	0.137
SupportDummy	0.377	0.485	0.000	0.000	1.000
Financial Performance					
OperatingRevenueGrowth	0.190	0.540	−0.217	0.107	0.560
OperatingProfitsGrowth	−0.198	3.311	−1.029	0.063	1.183
NetIncomeGrowth	−0.536	4.469	−1.021	0.060	1.126
SellExpensesGrowth	0.195	0.731	−0.342	0.086	0.649
AdminExpensesGrowth	0.172	0.462	−0.213	0.097	0.561
FinExpensesGrowth	0.094	2.956	−1.135	0.000	1.467
RDSpendingGrowth	0.322	1.028	−0.183	0.102	0.749
TotalAssetsGrowth	0.156	0.334	−0.091	0.088	0.437
CashGrowth	0.335	1.234	−0.447	0.007	1.176
AccountsReceivableGrowth	0.283	1.085	−0.354	0.064	0.820
FixedAssetsGrowth	0.233	0.792	−0.126	0.029	0.658
CAPEXGrowth	0.918	3.801	−0.637	0.000	2.118
TotalLiabGrowth	0.245	0.608	−0.175	0.091	0.742
STLoanGrowth	0.465	1.994	−0.632	0.006	1.336
LTLoanGrowth	0.921	4.812	−1.000	−0.002	1.890
ROA	0.032	0.079	−0.030	0.036	0.107
ROE	0.051	0.175	−0.048	0.069	0.183
Stock Performance					
Skewness	0.078	0.443	−0.442	0.065	0.608
Volatility	0.029	0.009	0.019	0.028	0.042
IdiosyncraticVolatility	0.024	0.008	0.015	0.023	0.034
AnnualReturn	0.138	0.590	−0.430	−0.004	0.896
AnnualVolume(Bil.)	31.727	41.367	4.556	17.999	71.153
MarketBeta	1.059	0.259	0.734	1.061	1.370

Table 3: **Policy Support and Firm-Level Outcomes**
This table reports estimates of the association between central government policy support and subsequent firm outcomes. $SupportStr_{i,t}$ is the number of supportive policies applicable to firm i's industry divided by the total number of supportive policies issued in year t. The dependent variable is the effective tax rate or the growth rate of the indicated outcome in year $t+1$. Firm controls include size, age, leverage, asset tangibility, political connections, operating revenue, and return on assets, as applicable. All specifications include firm and year fixed effects. Financial firms are excluded, and industry-clustered standard errors are reported in parentheses.

Panel A: Access to resources				
	Effective tax rate	Cash growth	Short-term loan growth	Long-term loan growth
	(1)	(2)	(3)	(4)
SupportStr	-0.051^{***}	0.724^{*}	0.818^{**}	0.368
	(0.019)	(0.366)	(0.354)	(1.028)
Firm controls	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	38,066	38,226	29,975	20,742
Adjusted R^2	0.226	0.068	0.062	0.007
Panel B: Operations and investment				
	Operating revenue growth	Operating profit growth	Net income growth	R&D spending growth
	(1)	(2)	(3)	(4)
SupportStr	0.231^{**}	-1.222^{**}	-1.027^{*}	0.478^{***}
	(0.108)	(0.495)	(0.611)	(0.133)
Firm controls	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	38,197	32,544	33,926	28,446
Adjusted R^2	0.129	0.103	0.128	0.046

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table 4: **Downside-Risk Changes Around Supportive-Policy Announcements**
This table compares downside risk for small firms before and after supportive-policy announcements. Panel A reports first announcements; Panel B includes all supportive-policy announcements. Pre covers months −12 through −1 relative to the announcement month, and Post covers months 0 through 11. Each row uses the same matched firm-events in both windows, with at least 120 daily observations per firm-window. Change is Post minus Pre. All measures are oriented so that larger values indicate greater downside risk. NCSKEW and DUVOL are daily crash-risk measures based on demeaned raw returns. p-values use standard errors clustered by industry.

Downside-risk measure	Pre	Post	Change	p-value
Panel A: First announcements				
Maximum drawdown magnitude (%)	40.756	36.777	−3.979	0.114
5th-percentile loss magnitude (%)	4.905	4.086	−0.819*	0.053
Frequency of daily returns at or below -5% (p.p.)	4.931	3.274	−1.657*	0.050
Lower-partial volatility (annualized %)	31.826	27.952	−3.874*	0.067
NCSKEW (unitless)	-0.011	-0.212	−0.201***	0.001
DUVOL (unitless)	0.105	-0.054	−0.159***	< 0.001
Panel B: All announcements				
Maximum drawdown magnitude (%)	41.421	40.268	−1.153***	0.008
5th-percentile loss magnitude (%)	5.393	5.015	−0.377***	0.001
Frequency of daily returns at or below -5% (p.p.)	5.916	5.240	−0.676***	0.003
Lower-partial volatility (annualized %)	34.475	33.002	−1.473***	0.003
NCSKEW (unitless)	-0.018	-0.061	−0.042***	0.001
DUVOL (unitless)	0.060	0.017	−0.044***	< 0.001

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Table 5: **Abnormal Returns to Supported and Non-Supported Portfolios**
This table reports monthly CAPM, Carhart, and CH4 alphas for small supported, large supported, and small non-supported portfolios. At the beginning of each month, a firm enters a supported portfolio if its industry received at least one supportive State Council policy document during the preceding 12 months. The non-supported portfolio is the complement and includes firms whose industries received only non-supportive policy documents or no policy document during the same window. Each portfolio is held for one month and rebalanced monthly. Small and large firms are the bottom and top halves of the monthly eligible market-capitalization distribution. The final row in each panel reports the directly estimated small supported-minus-non-supported spread. Panel A reports equal-weighted portfolios, and Panel B reports value-weighted portfolios. All portfolios use the baseline sample screens. The sample period is from June 2009 to December 2021. Alphas are monthly decimals, and standard errors are reported in parentheses. Appendix Tables J1 and J2 report full factor loadings for supported portfolios; Table J3 reports the small-firm spread.

	CAPM	Carhart	CH4
Panel A: Equal-weighted portfolios			
Small supported	0.012**	0.006***	0.006***
	(0.005)	(0.001)	(0.001)
Large supported	0.003	0.000	0.005**
	(0.003)	(0.002)	(0.002)
Small non-supported	0.009**	0.004***	0.000
	(0.004)	(0.001)	(0.001)
Small supported − non-supported	0.003*	0.002*	0.006***
	(0.001)	(0.001)	(0.002)
Panel B: Value-weighted portfolios			
Small supported	0.010**	0.005***	0.006***
	(0.004)	(0.001)	(0.001)
Large supported	0.001	0.002	0.003
	(0.002)	(0.002)	(0.002)
Small non-supported	0.007*	0.002**	0.000
	(0.004)	(0.001)	(0.001)
Small supported − non-supported	0.003*	0.002*	0.006***
	(0.002)	(0.001)	(0.002)

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Table 6: **Candidate Adverse States and the Policy-Portfolio Return Spread**
This table reports separate regressions of the value-weighted supported-minus-non-supported small-firm return spread on the CH4 factors and each adverse-state indicator. Panels A and B use adverse-quartile indicators. In Panel A, adverse states are the bottom quartile of monthly market excess returns and the top quartile of monthly market volatility or either economic policy uncertainty (EPU) index. Panel B uses four indicators of weak household and demand conditions defined in Appendix G.3. Consumer confidence and manufacturing PMI vary monthly. Each quarterly PBOC income- or employment-perception indicator is assigned to all three months of the corresponding quarter. Panel C uses three-month windows around four China-specific and three IMF-motivated policy-regime episodes. State coefficients are factor-adjusted adverse-minus-normal monthly return differences. Coefficients and standard errors are in monthly decimal returns. All panels use one-sided lower-tail tests of whether the state coefficient is negative. Portfolios are formed at the end of $t-1$ using support during the preceding 12 months.

Candidate adverse state	State coefficient	S.E.	p
Panel A: Conventional financial-market states			
Bottom-quartile market return	0.005	0.005	0.823
Top-quartile market volatility	0.010	0.004	0.996
Top-quartile Mainland EPU	−0.002	0.004	0.266
Top-quartile SCMP EPU	−0.005	0.004	0.110
Panel B: China household and demand states			
Bottom-quartile consumer confidence	0.005	0.004	0.919
Bottom-quartile manufacturing PMI	0.006	0.004	0.942
Bottom-quartile PBOC income perception	0.005	0.003	0.932
Bottom-quartile PBOC employment perception	0.007	0.003	0.972
Panel C: Externally dated policy-regime episodes			
Four China-specific policy-direction episodes	-0.013^{**}	0.006	0.012
Three IMF-motivated policy-regime episodes	-0.010^{*}	0.007	0.079

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table 7: **CH4 Alphas by Policy-Support Exposure**
This table reports monthly CH4 excess-return alphas for supported small firms across two measures of policy-support exposure. Portfolio formation in all panels follows the baseline construction in Table 5. Panel A compares SOEs with non-SOEs. Panel B compares firms whose industry is, or is not, a key-supported industry in the active FYP of the firm's registered province; firms without a detailed industry match are excluded. Long-only portfolio returns are measured in excess of the monthly risk-free rate. Difference rows estimate CH4 alphas directly on zero-cost monthly spreads. Entries are monthly decimals, with standard errors in parentheses.

Portfolio or contrast	Equal-weighted	Value-weighted
Panel A: Ownership exposure		
Supported small non-SOE	0.009***	0.009***
	(0.002)	(0.002)
Supported small SOE	0.001	0.001
	(0.002)	(0.002)
Non-SOE − SOE	0.008***	0.009***
	(0.003)	(0.003)
Panel B: Local Five-Year-Plan priority exposure		
Local-FYP priority match	0.008***	0.008***
	(0.002)	(0.002)
No local-FYP priority match	0.001	0.001
	(0.003)	(0.003)
Local priority − no local priority	0.007**	0.007**
	(0.003)	(0.003)

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table 8: **Robustness of Small-Firm Policy-Portfolio CH4 Alphas**
This table shows whether the monthly CH4 alpha of the supported small-firm portfolio persists under alternative samples, policy-support classifications, and systematic-risk controls. The first analysis in Panel A excludes technology-sector firms, and the second analysis removes the top decile of best-performing industries. Panel B augments the CH4 model with the SCMP-based EPU factor, the RDGD-based EPU factor, or the R&D factor; Appendix I describes their construction. Panel C replaces the baseline title-keyword classification with manual classification, GPT-based classification, or the intersection of the title-keyword, manual, and GPT classifications. Entries are monthly alphas in decimal returns. OLS standard errors are reported in parentheses. Factor loadings and goodness-of-fit statistics are omitted.

	Monthly CH4 alpha	
Robustness specification	Equal-weighted	Value-weighted
Panel A: Excluding high-return industries		
Excluding the technology sector	0.006***	0.005***
	(0.002)	(0.002)
Excluding ex-post top-decile return industries	0.003**	0.003**
	(0.002)	(0.002)
Panel B: Additional systematic-risk controls		
CH4 + SCMP-based EPU factor	0.006***	0.006***
	(0.001)	(0.001)
CH4 + RDGD-based EPU factor	0.006***	0.006***
	(0.001)	(0.001)
CH4 + R&D factor	0.006***	0.006***
	(0.001)	(0.001)
Panel C: Alternative policy-support classifications		
Manual policy-support classification	0.007***	0.007***
	(0.002)	(0.002)
GPT-based policy-support classification	0.007***	0.007***
	(0.002)	(0.002)
Intersection of title-keyword, manual, and GPT classifications	0.007***	0.007***
	(0.002)	(0.002)

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table 9: **Policy-Classification Robustness of the Small-Firm CH4 Alpha**
This table reports small-firm CH4 alphas under alternative policy-support classifications and support-content categories. Documents and Links count the supportive documents and document–industry links underlying each classification row. Panel A reports the baseline 282-document, 787-link classification and two exclusions of cases triggered only by: three generic terms (*reform*, *plan*, or *strengthen*) or the three largest sole-trigger words (*strengthen*, *accelerate*, or *reform*). Panel B uses classifications from two variants of the DeepSeek model: V4-Pro and V4-Flash. Panel C uses DeepSeek V4-Pro to distinguish resource/cash-flow support, executable nonresource support, and vague guidance, ordered from highest to lowest support intensity. Alphas are monthly decimals, with standard errors in parentheses.

Policy-support definition	Documents	Links	Equal-weighted	Value-weighted
Panel A: Title-keyword sensitivity				
Baseline title-keyword rule	282	787	0.006***	0.006***
			(0.001)	(0.001)
Exclude reform, plan, and strengthen	229	643	0.006***	0.006***
			(0.001)	(0.001)
Exclude strengthen, accelerate, and reform	241	635	0.006***	0.006***
			(0.001)	(0.001)
Panel B: Repeated full-text AI classifications				
DeepSeek V4-Pro majority	272	539	0.006***	0.006***
			(0.001)	(0.001)
DeepSeek V4-Flash majority	296	617	0.005***	0.005***
			(0.002)	(0.002)
Panel C: V4-Pro majority support-content categories				
Resource/cash-flow support (intensity 3)	241	381	0.006***	0.006***
			(0.002)	(0.002)
Executable nonresource support (intensity 2)	31	53	0.000	−0.001
			(0.005)	(0.005)
Vague guidance (intensity 1)	38	95	0.005	0.005
			(0.005)	(0.005)

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Appendix A Policy-Document Measurement

A.1 Database Filters and Baseline Support Rule

The sample construction described in Section 3.1 begins with 5,762 records issued by the State Council or its General Office. Restricting publication dates to after March 28, 2008 leaves 2,663 records. The administrative-title filter retains 1,574 documents. Its 11 terms are 决定 (*jueding*, decision), 办法 (*banfa*, measures), 意见 (*yijian*, opinion), 措施 (*cuoshi*, measures), 方案 (*fang'an*, plan or program), 条例 (*tiaoli*, regulations), 细则 (*xize*, detailed rules), 规划 (*guihua*, plan), 规则 (*guize*, rules), 规定 (*guiding*, provisions), and 规范 (*guifan*, norms or standards). Industry mapping yields 505 documents. Removing 47 non-substantive records and 54 region-specific documents leaves 404 policies, corresponding to 1,074 distinct document–industry links in 60 CSRC industries.

In the baseline analysis, we use the following title-support keywords to identify 282 supportive documents and 787 supportive links. The keywords are: 保护 (*baohu*, protect), 促进 (*cujin*, promote), 发展 (*fazhan*, develop), 改革 (*gaige*, reform), 规划 (*guihua*, plan), 鼓励 (*guli*, encourage), 加快 (*jiakuai*, accelerate), 健全 (*jianquan*, improve), 加强 (*jiaqiang*, strengthen), 积极 (*jiji*, actively), 进一步 (*jinyibu*, further), 开展 (*kaizhan*, carry out), 深化 (*shenhua*, deepen), 深入 (*shenru*, in depth), 提升 (*tisheng*, enhance), 统一 (*tongyi*, unify), 推动 (*tuidong*, push forward), 推进 (*tuijin*, advance), 行动 (*xingdong*, action), 增长 (*zengzhang*, growth), 支持 (*zhichi*, support), and 振兴 (*zhenxing*, revitalize).

A.2 Theme-to-Industry Mapping

The State Council database assigns documents to 95 official themes. We map 34 themes with an identifiable industry exposure to one or more industries in the CSRC's 2012 Industry Classification of Listed Companies. Table 1 reports the complete mapping. Themes without a natural listed-industry counterpart include many policies concerning population and family planning, religion, cultural relics, public security, administrative place names, ethnic affairs, and women and children.

The technology and intellectual-property themes map to high-technology industries defined by the National Bureau of Statistics of China. We map a two-digit CSRC industry to high technology when more than half of the corresponding subcategories fall within the

national high-technology classification.

A.3 Manual and GPT Classifications

For robustness, we use manual reading and GPT to identify policy support from the content of policy documents, providing alternatives to the baseline title-keyword rule. In the manual classification, we read each document and classify it as supportive or non-supportive. This document-level classification applies to all industries mapped to the document. Manual reading identifies 222 supportive documents, corresponding to 751 document–industry links.

For the GPT classification, we provide GPT-3.5 with the policy text and the name of a mapped industry. The model assesses whether the policy has a positive, negative, or uncertain effect on the market value of an average firm in that industry. Unlike the document-level manual classification, GPT evaluates each document–industry link separately.[23] The classification task is stated as follows:

> You are a financial analyst evaluating the impact of a new policy by the State Council of China on the market value of an average Chinese company in the given industry. Judge whether this policy will have a positive, negative, or uncertain effect on the market value of the firm. Answer only with a number: 1 if positive, -1 if negative, and 0 if uncertain.

GPT classifies 595 links as positive, 469 as uncertain, and 10 as negative. We treat positive assessments as supportive and pool uncertain and negative assessments as non-supportive.

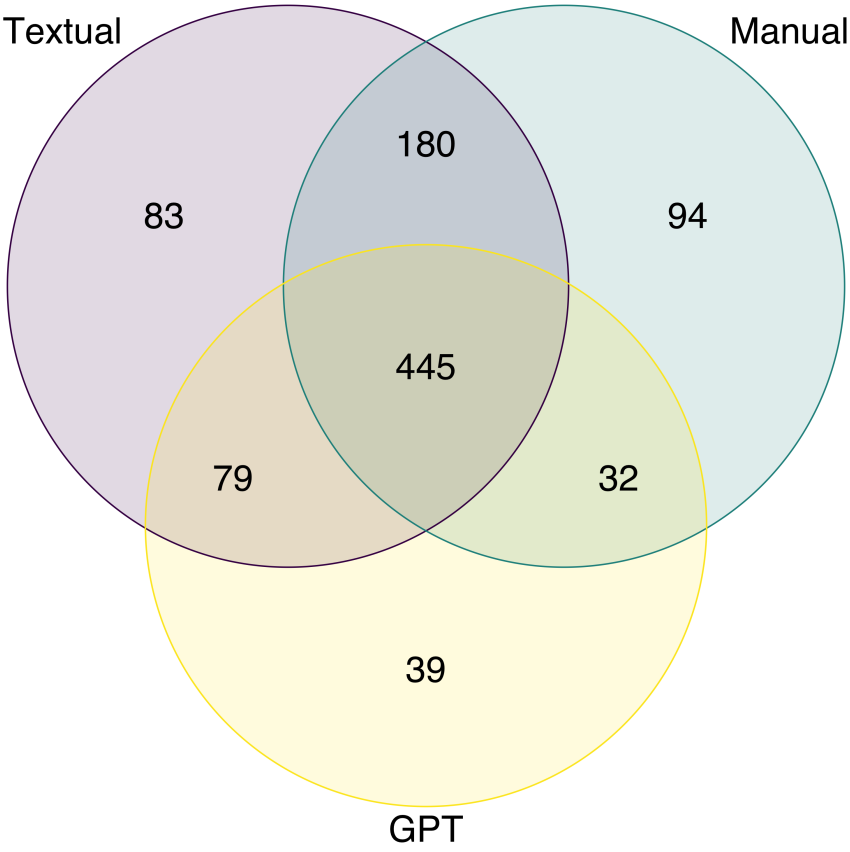
Figure A1 compares the title-keyword, manual, and GPT classifications across the same 1,074 links from 404 policies. These methods identify 787 (73.3%), 751 (69.9%), and 595 (55.4%) links as supportive, respectively. Although GPT classifies fewer links as supportive than manual reading, the two methods overlap substantially: 477 of the 595 GPT-supportive links (80.2%) are also classified as supportive manually. All three methods identify 445 links as supportive.[24]

[23] The model specification is `gpt-3.5-turbo-0301`, with temperature set to 1. For 87 links from 23 documents, the input contains the first 12,000 characters of the policy text, followed by the industry's name.

[24] Counting agreement on both supportive and non-supportive classifications, overall agreement is 73.2% between the title-keyword and manual methods, 68.9% between the title-keyword and GPT methods, and 63.5% between the manual and GPT methods. These percentages differ from the share of GPT-supportive links also classified as supportive manually.

The return results are also consistent across the manual and GPT measures. As reported in Section 7.4, portfolios formed using either classification earn positive and statistically significant CH4 alphas (Table 8, Panel C). Thus, the positive abnormal-return finding is not specific to the baseline title-keyword rule.

Figure A1: **Alternative Policy-Support Classifications**
This figure compares supportive document–industry links identified by the title-keyword rule (Textual), manual reading (Manual), and GPT. Manual classifications are assigned at the document level and applied to its mapped industries; GPT evaluates each document–industry link separately. The sample contains 404 documents and 1,074 links. The circles contain 787, 751, and 595 supportive links, respectively. Manual and GPT classifications share 477 supportive links, of which 445 are also identified by the title-keyword rule.



A.4 DeepSeek Classification of Policy Support and Its Content

We additionally use DeepSeek to classify policies and distinguish favorable policy language from concrete measures that may benefit firms. This classification complements the manual and GPT measures by identifying how a document supports each mapped industry. We provide the full Chinese text, title, document type, official theme, issuing body, publication date, and industry codes and names. The model evaluates the policy's direction for an average listed firm in each industry and extracts the measures underlying that assessment. Neither stock returns nor firm outcomes are supplied to the model.

We apply the same Chinese-language instructions to DeepSeek V4-Pro and V4-Flash, repeating the exercise three times for each model. Every application covers all 404 documents and 1,074 document–industry links. The instructions are reproduced at the end of this appendix.[25]

For each industry, DeepSeek reports whether the document is beneficial, adverse, or neutral/uncertain, with a supporting quotation, confidence assessment, and short explanation. It also counts distinct executable measures and describes up to four principal measures. Each description identifies the policy channel, whether the measure is executable or provides resources, and any explicit budget, amount, funding source, responsible agency, eligibility condition, or deadline. The instructions require these details to come from the document; absent details are left blank. The count may exceed the four measures described and does not determine the intensity classification.

An executable measure specifies an implementation instrument, obligation, eligibility rule, approval procedure, standard, responsible body, or other operational arrangement. General encouragement or a statement of policy goals does not suffice. Resource support requires a favorable cash-flow or resource provision through one of eight channels: fiscal transfers, credit, tax, public procurement, pricing, market access, mandated demand, or resource allocation. General regulation, planning, and technical standards do not qualify as resource support on their own. We count an extracted measure toward executable or resource support only when its supporting quotation appears verbatim in the source document.

Beneficial direction gives support $+1$, adverse direction gives -1, and neutral/uncertain

[25] The classifications were obtained using `deepseek-v4-pro` and `deepseek-v4-flash`, yielding 2,424 document responses and 6,444 link-level assessments across the six applications. Thinking mode was disabled. The request code sets temperature to zero and supplies no random seed.

direction gives zero. Links without favorable support receive intensity zero. Supportive links receive intensity three if they contain a qualifying resource measure, intensity two if they contain an executable measure but no qualifying resource measure, and intensity one otherwise. These categories correspond to resource/cash-flow support, executable nonresource support, and vague guidance. Cash-flow support requires both supportive direction and a qualifying resource measure; vague guidance requires supportive direction and intensity one.

For each model, we assign the majority classification across its three applications separately for support, cash-flow support, vague guidance, and intensity. A three-way disagreement leaves that field unclassified. This occurs only for intensity, affecting ten Pro links and twelve Flash links. A supportive link with unresolved intensity remains in the support comparison but cannot enter an intensity category.[26]

Panel A of Table A1 summarizes the four support classifications. The Pro classification overlaps substantially with manual coding: 428 of its 539 supportive links (79.4%) are also manually supportive. Table 9 reports the return results for both DeepSeek variants.

Panel B shows that resource support is the most common form of support. Of the 404 documents, 259 (64.1%) contain executable nonresource or resource support for at least one mapped industry. Because a document can support different industries in different ways, the document-level categories overlap.

[26]The measure count is the rounded median across applications, capped at ten. A policy channel is retained when at least two applications identify it.

Table A1: **Policy-Support Classifications and Content**
This table reports classification counts for 404 documents and 1,074 document–industry links. Panel A compares four support measures. Panel B reports three categories of supportive content under the V4-Pro majority classification, based on three applications. Documents can cover several industries and support categories, so document counts need not sum across rows.

Classification	Links	Documents
Panel A: Support definitions		
Title-keyword support	787	282
Manual document coding	751	222
GPT positive label	595	234
DeepSeek V4-Pro majority support	539	272
Panel B: V4-Pro majority intensity		
Vague guidance (intensity 1)	95	38
Executable nonresource support (intensity 2)	53	31
Resource/cash-flow support (intensity 3)	381	241

A.5 Consistency of the DeepSeek Classifications

Table A2 shows high within-model repeatability and lower cross-model agreement. The Pro and Flash majority labels agree on all four classifications for 824 of 1,074 links (76.7%).

Table A2: **Repeatability and Cross-Model Agreement**
This table reports agreement rates across repeated applications of the same model and across Pro and Flash majority classifications for all 1,074 document–industry links. The V4-Pro and V4-Flash columns report average pairwise agreement across three applications. The cross-model column reports agreement between Pro and Flash majority labels, treating missing fields as disagreements.

	V4-Pro	V4-Flash	Cross-model
Support agreement (%)	94.9	95.8	87.4
Cash-flow agreement (%)	94.2	94.0	87.2
Vague-guidance agreement (%)	95.0	95.4	90.5
Intensity agreement (%)	90.6	90.1	78.5

A.6 Chinese-Language Instructions for Policy Classification

For reproducibility, we reproduce the Chinese-language instructions given to both DeepSeek models. They ask the model to assess policy direction for each mapped industry, identify concrete measures, and quote the supporting text. The system instruction defines the task; the user instruction specifies the classification rules and the structured response format (JSON). The example illustrates that format rather than an observed policy or model response. The original instructions and example are reproduced without substantive changes.

System instruction. 你正在为一篇实证金融论文提取中国国务院产业政策文件中的可核验事实。只返回合法JSON，不要使用markdown，不要增加schema 之外的字段。请针对列出的每一个CSRC 行业分别提取，判断对象是该行业内一家平均中国上市公司。必须依据政策全文，不得仅依据标题、文件类型或一般常识。所有evidence_phrase 必须是政策原文中可以逐字检索到的连续中文短语。

User instruction. 请按照下面的JSON schema 返回一个JSON object，并为每个目标行业返回且只返回一行。

提取规则:

1. direction 只能是beneficial、adverse、neutral_uncertain。beneficial 表示文件对该行业平均上市公司有明确有利方向；adverse 表示明确不利；中性、无法判断、并非针对该行业时用neutral_uncertain。

2. direction_evidence_phrase 必须是支持direction 判断的原文连续短语，即使direction 为neutral_uncertain 也要给出最相关原文。

3. actionable_measure_count 是全文中与该目标行业直接相关的不同具体可执行措施数量。不要把口号、背景介绍、目标愿景或重复表述计入。

4. measures 最多列出四项最重要的具体措施；没有具体措施时返回空数组。每项措施必须有一条逐字原文evidence_phrase。

5. channel 只能从以下英文标签中选择：fiscal, credit, tax, procurement, pricing, market_access, mandated_demand, resource_allocation, regulation, planning, other。

6. binding_or_executable=1 仅限有明确执行工具、义务、资格、审批、标准、实施主体或可操作安排的措施；纯粹鼓励、推动、加强、规划方向填0。

7. cashflow_or_resource=1 仅限通过财政、信贷、税收、政府采购、价格、市场准入、强制需求或明确资源配置影响企业现金流的有利措施。一般监管、规划或技术标准本身不得填1。

8. explicit_budget=1 仅限原文明确写出财政预算、专项资金或金额。不得推断隐含预算。explicit_amount 必须逐字记录原文金额；没有则为空字符串。

9. funding_source、responsible_agency、eligibility_condition、deadline 只记录原文明示内容，没有则为空字符串。

10. 不要因为标题出现发展、改革、规划、推进、加强、意见等词就判断为支持或高强度。监管合规成本也不能被误判为支持。

11. confidence 在0 到1 之间。short_rationale 使用中文，最多两句话。

Illustrative response format (JSON).

```
{
  "pub_url": "...",
  "document_summary": "中文一句话摘要",
  "industry_annotations": [
    {
      "pub_url": "...",
      "doc_title": "...",
      "ind_code": 44,
      "industry_name": "Electric power, heat production and supply",
      "direction": "beneficial",
      "direction_evidence_phrase": "支持符合条件的项目建设",
      "actionable_measure_count": 2,
      "measures": [
        {
          "channel": "credit",
          "binding_or_executable": 1,
          "cashflow_or_resource": 1,
          "explicit_budget": 0,
          "explicit_amount": "",
          "funding_source": "政策性银行",
          "responsible_agency": "有关金融机构",
          "eligibility_condition": "符合条件的项目",
          "deadline": "",
          "evidence_phrase": "加大对符合条件项目的信贷支持"
        }
      ],
      "confidence": 0.85,
      "short_rationale": "中文简短理由"
    }
  ]
}
```

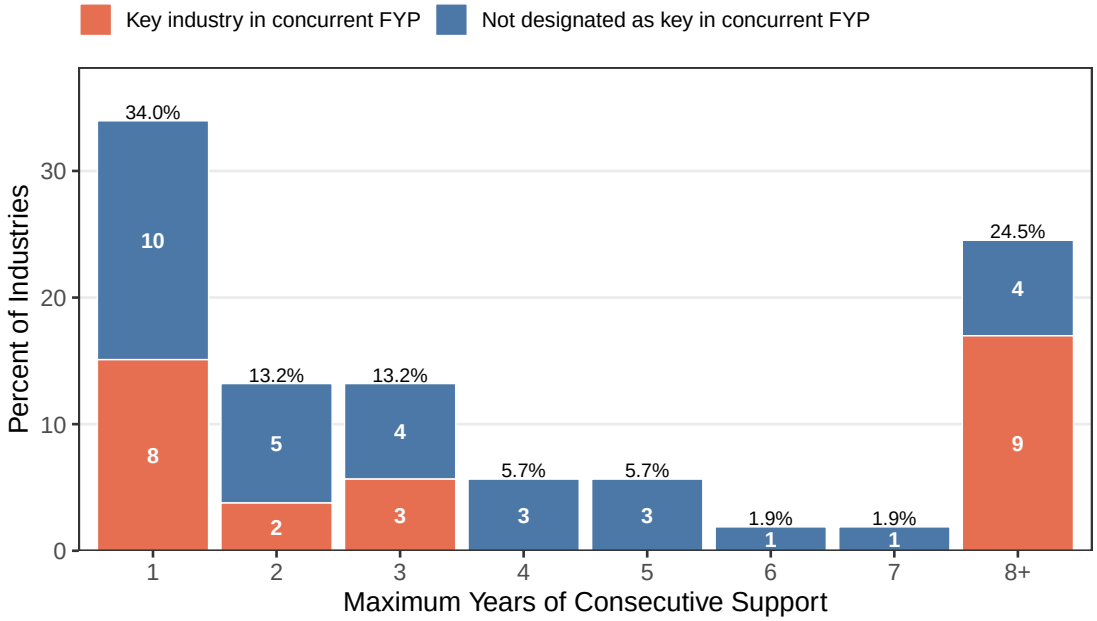
Document information supplied to the model. Each classification request follows these instructions with the document's URL, title, type, official theme, issuing body, and publication date, in that order. It then lists all mapped industries, giving each industry's code and English and Chinese names, and concludes with the complete Chinese policy text. These document-specific inputs accompany the same instructions for both models and all repeated applications.

Appendix B Five-Year Plans and the Duration of Policy Support

This appendix examines whether the long support spells in Panel B of Figure 3 are concentrated in industries designated as central Five-Year Plan (FYP) priorities. Figure B1 separates the same duration distribution by priority status. The analysis describes the association between plan priorities and realized support duration. The distribution remains bimodal among industries designated as central Five-Year Plan priorities.

We obtain central-plan priorities from the CNRDS Industrial Policy Research Database. An industry is classified as a key priority if at least one matched record designates it as key-supported in the central FYP cycle covering the first calendar year of its longest observed support spell. We map CSRC-2012 industries to CNRDS's CSRC-2001 codes using the old and new codes in CSMAR company profiles.

Figure B1: **Maximum Support Duration by Central Five-Year Plan Priority**
This figure replicates Panel B of Figure 3 but distinguishes industries by their central FYP priority status. Each observation is one of the 53 two-digit CSRC industries with at least one supported calendar year during 2008–2021. Duration is the longest observed run of consecutive years with at least one supportive State Council document. Orange identifies key priorities in the central FYP cycle covering the first calendar year of the selected spell, and blue identifies the remaining industries. CNRDS plan classifications are matched through CSMAR's CSRC-2001/2012 crosswalk.



Appendix C Policy Support and Firm-Level Outcomes

This appendix provides the full firm-level financial statement results presented in Section 4.1.

We use the following fixed-effects regression as our baseline specification:

$$Y_{i,t+1} = \beta_1 SupportStr_{i,t} + \beta_{\mathbf{2}}' \mathbf{X_{i,t}} + a_i + a_t + u_{i,t} \tag{12}$$

where $Y_{i,t+1}$ denotes a measure of firm i's performance in year $t+1$; the outcomes include both balance-sheet and income-statement items. $SupportStr_{i,t}$ measures the support strength for firm i in year t and is calculated as the number of supportive policies applicable to firm i divided by the number of total supportive policies announced in year t. The unconditional median is zero and the mean is 4.5%. Thus, on average, nearly 5% of the supportive policies issued in a given year apply to each firm's industry.

$\mathbf{X_{i,t}}$ denotes the vector of firm-level controls, which, depending on the specification, includes size (the logarithm of total assets), age (the logarithm of the number of months since IPO), leverage, tangibility (fixed assets divided by total assets), political connection (the logarithm of the number of board members who are or were central government officials), the logarithm of operating revenue, and ROA. a_i and a_t are firm and year fixed effects, respectively. We cluster standard errors at the industry level.

The summary statistics for this analysis are reported in Table 2 in the main body of the paper. Following standard practice, we exclude financial firms when analyzing performance.

C.1 Policy and Balance Sheet Health

Central government industrial policy is implemented by local government officials who have control or influence over local assets, including local banks (Li et al., 2008; Liu, 2023). Strong policy support should thus be directly related to a firm's cash position and its ability to obtain loans.

Table C1 reports the β_1 coefficients for the relationship between industrial policy support strength and the growth in various balance sheet items, including total assets, cash, accounts receivable, fixed assets, total liabilities, short-term loans, and long-term loans.

Stronger policy support is associated with growth in both assets and liabilities. On the

asset side, policy support is positively related to the growth in total assets. Cash holdings and accounts receivable increase, but investment in tangible assets is not significantly affected, as measured by fixed asset growth. Total liabilities increase and this is driven by the growth of short-term, rather than long-term, loans.

The point estimates are economically significant. A one-standard-deviation increase in support strength (0.070) is associated with 5.07 percentage points (p.p.) higher cash growth, 3.01 p.p. higher accounts-receivable growth, and 5.73 p.p. higher short-term-loan growth.

Stronger policy support is associated with easier access to credit, higher cash holdings, and more trade credit (i.e., accounts receivable).

Table C1: **Policy Support and Firm-Level Growth: Balance Sheet**
This table presents coefficients from the following fixed-effects regressions: $Y_{i,t+1} = \beta_1 SupportStr_{i,t} + \beta_2' \mathbf{X_{i,t}} + a_i + a_t + u_{i,t}$. $Y_{i,t+1}$ denotes measures of firm i's growth in assets and liabilities in year $t+1$, including total assets ("TA"), cash, accounts receivable ("AR"), fixed assets ("FA"), total liabilities ("TL"), short-term loans ("ST Loan"), and long-term loans ("LT Loan"). $SupportStr_{i,t}$ measures the policy support strength for firm i in year t and is calculated as the number of announced supportive policies for the corresponding industry divided by the number of total supportive policies announced in year t. $\mathbf{X_{i,t}}$ is a vector of firm-level controls including size (logarithm of total assets), age (logarithm of the number of months since IPO), leverage, tangibility (fixed assets divided by total assets), political connection (logarithm of the number of board members who are or were central government officials), the logarithm of operating revenue, and ROA. a_i and a_t are firm and year fixed effects, respectively. Industry-clustered standard errors are reported in parentheses.

	Dependent variable:						
	TA Growth	Cash Growth	AR Growth	FA Growth	TL Growth	ST Loan Growth	LT Loan Growth
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
SupportStr	0.152**	0.724*	0.430**	−0.069	0.344***	0.818**	0.368
	(0.065)	(0.366)	(0.178)	(0.125)	(0.098)	(0.354)	(1.028)
FirmSize	−0.223***	−0.574***	−0.189***	−0.145***	−0.301***	−0.362***	−1.433***
	(0.011)	(0.033)	(0.037)	(0.022)	(0.015)	(0.044)	(0.157)
Age	0.041***	0.180***	0.041***	−0.034***	0.017**	−0.163***	0.419***
	(0.005)	(0.014)	(0.014)	(0.012)	(0.008)	(0.040)	(0.108)
Leverage	0.020	0.998***	−0.192*	−0.041	−1.073***	−2.382***	−2.810***
	(0.025)	(0.102)	(0.097)	(0.067)	(0.059)	(0.182)	(0.459)
Tangibility	−0.030	1.000***	−0.056	−2.349***	−0.179**	−0.154	−2.557***
	(0.040)	(0.117)	(0.106)	(0.134)	(0.069)	(0.172)	(0.544)
PoliticalConnection	0.012**	−0.010	0.017	0.019	0.011	0.006	0.062
	(0.006)	(0.016)	(0.019)	(0.015)	(0.012)	(0.044)	(0.112)
OperatingRevenue	0.035***	−0.013	−0.181***	−0.077***	0.007	0.007	0.377***
	(0.009)	(0.031)	(0.030)	(0.018)	(0.010)	(0.036)	(0.122)
ROA	0.750***	0.373	0.780***	0.556***	0.343***	0.799***	1.678**
	(0.054)	(0.235)	(0.158)	(0.108)	(0.078)	(0.258)	(0.663)
Firm + Year FE	Y	Y	Y	Y	Y	Y	Y
Industry Clustering	Y	Y	Y	Y	Y	Y	Y
Observations	38,252	38,226	37,873	38,220	38,227	29,975	20,742
R^2	0.247	0.167	0.140	0.220	0.233	0.175	0.149
Adjusted R^2	0.158	0.068	0.037	0.127	0.142	0.062	0.007

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

C.2 Policy and Income Statement Analysis

Turning to the income statement, we find that firms pay lower taxes and, consistent with the increase in accounts receivable documented above, generate more operating revenue. Firms spend substantially more on R&D, but operating profits and net income fall.

Table C2 reports the relationships between policy support and effective tax rates and between policy support and growth in revenue and expenses, including operating revenue, operating profits, net income, selling expenses, administrative expenses, financial expenses, and R&D spending.

Support strength is related to a decrease in the effective tax rate. A one-standard-deviation increase in support strength is related to a 0.36 p.p. reduction in the ETR. This is consistent with the literature, which documents that many forms of business tax fall under the discretion of the local government (Cai, Fang, and Xu, 2011).

Support strength is positively related to operating revenue growth the following year. However, the increase does not translate into higher operating profits or net income. Operating profits and net income decrease, likely because operating expenses grow more rapidly. We therefore examine how policy support relates to different types of expenses. The results in columns 5 through 7 indicate that firms face increased selling expenses after receiving supportive policies, while the growth in administrative expenses becomes slower. Support strength is not related to the growth in financial expenses.

Increased expenses may also show up in R&D spending, which includes both capitalized R&D expenditures and R&D expenses. Column 8 shows that policy support is positively related to growth in R&D spending. A one-standard-deviation increase in support strength is associated with a 3.35 p.p. increase in R&D spending growth.

Stronger policy support is associated with more operating activity, lower effective tax rates, slower growth in administrative expenses, and greater R&D investment.

Table C2: **Policy Support and Firm-Level Growth: Income Statement**
This table presents coefficients from the following fixed-effects regressions: $Y_{i,t+1} = \beta_1 SupportStr_{i,t} + \beta_{\mathbf{2}}' \mathbf{X_{i,t}} + a_i + a_t + u_{i,t}$. $Y_{i,t+1}$ denotes measures of firm i's effective tax rate and growth in revenue and expenses in year $t+1$, including operating revenue ("OR"), operating profits ("OP"), net income ("NI"), selling expenses ("Sell Exp"), administrative expenses ("Admin Exp"), financial expenses ("Fin Exp"), and R&D spending ("R&D Spend"). $SupportStr_{i,t}$ measures the policy support strength for firm i in year t and is calculated as the number of announced supportive policies for the corresponding industry divided by the number of total supportive policies announced in year t. $\mathbf{X_{i,t}}$ is a vector of firm-level controls including size (logarithm of total assets), age (logarithm of the number of months since IPO), leverage, tangibility (fixed assets divided by total assets), political connection (logarithm of the number of board members who are or were central government officials), the logarithm of operating revenue, and ROA. a_i and a_t are firm and year fixed effects, respectively. Industry-clustered standard errors are reported in parentheses.

	Dependent variable:							
	ETR	OR Growth	OP Growth	NI Growth	Sell Exp Growth	Admin Exp Growth	Fin Exp Growth	R&D Spend Growth
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SupportStr	-0.051^{***}	0.231^{**}	-1.222^{**}	-1.027^{*}	0.286^{**}	-0.284^{***}	0.284	0.478^{***}
	(0.019)	(0.108)	(0.495)	(0.611)	(0.134)	(0.090)	(0.339)	(0.133)
FirmSize	-0.002	0.151^{***}	-0.641^{***}	-0.802^{***}	0.050^{***}	-0.008	0.250^{***}	0.006
	(0.005)	(0.024)	(0.145)	(0.170)	(0.018)	(0.011)	(0.074)	(0.050)
Age	-0.001	0.051^{***}	-0.009	0.031	0.019^{**}	0.008	0.002	0.125^{***}
	(0.002)	(0.008)	(0.037)	(0.052)	(0.010)	(0.006)	(0.052)	(0.017)
Leverage	-0.028^{***}	0.151^{***}	0.075	-0.852^{**}	-0.031	-0.044^{*}	0.505^{**}	0.177^{**}
	(0.009)	(0.057)	(0.311)	(0.372)	(0.066)	(0.025)	(0.209)	(0.087)
Tangibility	-0.025	0.109	0.718^{*}	0.148	-0.119	-0.022	-0.088	-0.369^{*}
	(0.015)	(0.070)	(0.401)	(0.539)	(0.074)	(0.037)	(0.232)	(0.195)
PoliticalConnection	-0.003	0.013	-0.066	0.014	-0.005	0.000	0.052	0.034
	(0.002)	(0.008)	(0.068)	(0.069)	(0.011)	(0.005)	(0.042)	(0.021)
OperatingRevenue	0.016^{***}	-0.372^{***}	0.059	0.244^{*}	-0.220^{***}	-0.091^{***}	-0.100^{*}	-0.291^{***}
	(0.003)	(0.031)	(0.159)	(0.142)	(0.019)	(0.009)	(0.059)	(0.038)
ROA	0.030^{*}	0.755^{***}	8.924^{***}	16.168^{***}	0.933^{***}	0.712^{***}	0.647^{**}	1.354^{***}
	(0.018)	(0.077)	(1.339)	(1.597)	(0.113)	(0.066)	(0.303)	(0.172)
Firm + Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Industry Clustering	Y	Y	Y	Y	Y	Y	Y	Y
Observations	38,066	38,197	32,544	33,926	37,328	38,252	38,250	28,446
R^2	0.308	0.221	0.214	0.232	0.185	0.204	0.124	0.172
Adjusted R^2	0.226	0.129	0.103	0.128	0.086	0.110	0.021	0.046

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Appendix D Policy Actionability and Realized Benefits

This appendix extends the analysis in Section 4.1 and examines whether supportive policies with concrete fiscal, financial, or other resource provisions are associated with greater subsequent benefits for targeted firms.

We use the three-run DeepSeek V4-Pro majority classification described in Appendix A. We pool vague guidance and executable nonresource measures as *nonresource support* and compare them with *resource-oriented support*.

We restrict the analysis to industry-years with resolved support, covering 9,686 eligible firm-years before outcome-specific missing-data restrictions. Outcomes and firm controls are winsorized at the 1st and 99th percentiles.

For industry j in year t, let $R_{j,t}$ count distinct resource-oriented documents. Total support strength, $S_{j,t}^{Total}$, is the industry's resolved supportive-document count divided by the annual national count of unique V4-Pro-supportive documents. For firm i in industry j, we estimate

$$Y_{i,t+1} = \beta \mathbf{1}\{R_{j,t} > 0\} + \delta S_{j,t}^{Total} + \boldsymbol{\gamma}' \mathbf{X}_{i,t} + a_i + a_t + u_{i,t}. \tag{13}$$

The indicator identifies years with at least one resource-oriented document, including years with both types of support. The coefficient β compares the two groups at the same total support strength. Policy support and firm controls are measured in year t, and outcomes are measured in year $t+1$; the table note reports the fixed effects and clustering convention.

Table D1 reports an effective tax rate lower by 1.09 percentage points and short-term-loan growth higher by 30.1 percentage points for resource-present firm-years. The cash-growth and operating-profit-growth differences are positive but only marginally significant. No other displayed resource-presence coefficient is significant. The tax and borrowing results are consistent with fiscal and financial benefits for targeted firms.

Table D1: **Resource-Oriented versus Nonresource Policy Support and Firm Outcomes** This table compares resource-present with nonresource-only supported firm-years, controlling for total support strength. Nonresource support pools vague guidance and executable measures without qualifying resource provisions. Resource presence requires at least one resource-oriented document, including years with both types of support. Total support strength is the industry's supportive-document count divided by the annual national count of unique V4-Pro-supportive documents. Following Tables C1 and C2, policy support and controls are measured in t, and outcomes are measured in $t+1$. Controls comprise log assets, log listing age, leverage, tangibility, political connections, log operating revenue, and ROA. All models include firm and year fixed effects; parentheses report standard errors clustered by industry.

Access to resources				
	Effective tax rate	Cash growth	Short-term loan growth	Long-term loan growth
	(1)	(2)	(3)	(4)
Resource-oriented support	−0.011***	0.071*	0.301***	0.108
	(0.003)	(0.035)	(0.072)	(0.270)
Total SupportStr	−0.006	0.885**	1.107	−1.876
	(0.034)	(0.423)	(0.846)	(1.983)
Observations	7,682	8,973	6,646	4,183
R^2	0.500	0.231	0.264	0.225
Operations, investment, and profitability				
	Operating revenue growth	Operating profit growth	ROA	R&D spending growth
	(1)	(2)	(3)	(4)
Resource-oriented support	−0.004	0.134*	0.003	−0.055
	(0.017)	(0.073)	(0.002)	(0.041)
Total SupportStr	0.460*	−0.693	−0.003	1.308***
	(0.238)	(1.207)	(0.034)	(0.476)
Observations	8,976	7,602	8,984	7,168
R^2	0.321	0.305	0.478	0.264
Firm controls dated t	Yes	Yes	Yes	Yes
Both policy terms entered jointly	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Appendix E Policy Retrenchment and Continuation

E.1 Explicit Policy Retrenchment

The main analysis measures the continuation of support through subsequent State Council documents because the government rarely announces that an industrial-policy priority has failed or formally ended. Most changes in support must therefore be inferred from the appearance, content, and absence of later directives. A small number of documents in the sample contain language that more clearly reverses, restricts, or retrenches an earlier policy. These cases provide useful illustrations of realized policy change but are too rare and heterogeneous to support a separate portfolio test.

This appendix describes each explicit retrenchment case, the earlier policy objective, the later change in direction, and the affected industries. The discussion distinguishes an actual reversal of support from narrower adjustments to eligibility, subsidy intensity, licensing, capacity, or implementation rules. The cases complement the duration analysis by showing why a missing renewal generally provides a more useful empirical measure of declining support than an explicit termination announcement. Table E1 documents three explicit reductions in policy-support instruments: photovoltaic tariffs, wheat minimum purchase prices, and new-energy-vehicle purchase subsidies. A case requires an earlier supportive State Council or General Office document, a later action in the same policy program, an explicit reduction in the relevant support instrument, and original-language evidence from both documents. We searched later departmental policy documents for explicit reductions in support, using AI-assisted screening followed by verification against the original policy texts. Among the 787 supportive document–industry links in our sample, this procedure identified only three clear cases of explicit retrenchment.

Table E1: **Three Documented Reductions in Policy-Support Instruments**
This table pairs earlier supportive policies with three subsequent reductions in support instruments. Dates are signing dates. English titles are the authors' translations and link to the original documents.

Earlier supportive policy	**Later retrenchment policy**
Opinions on Promoting the Healthy Development of the Photovoltaic Industry July 4, 2013	*Notice on Matters Related to Photovoltaic Power Generation in 2018* May 31, 2018
Opinions on Improving Support Policies to Promote Sustained Growth in Farmers' Income November 24, 2016	*Notice Announcing the Minimum Purchase Price for Wheat in 2019* November 15, 2018
Guiding Opinions on Accelerating the Promotion and Application of New-Energy Vehicles July 14, 2014	*Notice on Further Improving Fiscal Subsidy Policies for the Promotion and Application of New-Energy Vehicles* March 26, 2019

Solar as an implementation-risk illustration. The 2013 solar directive supported deployment through tariffs, subsidies, finance, land, and grid access, while allowing tariff reductions as costs fell. Firms nevertheless committed capital before the later terms of support were known. On January 4, 2018, LONGi signed an agreement for a RMB 1.95 billion, 5-GW module project in Chuzhou.[27] The May 31 notice then curtailed new subsidized deployment and lowered relevant tariffs. LONGi's 2018 annual report attributes a reduction in its production plan to weaker domestic demand after the policy change. It also reports a 28.24% decline in net profit attributable to shareholders.[28]

[27]LONGi, *LONGi Solar Plans to Invest RMB 1.95 Billion to Build a 5GW Module Plant in China*, January 10, 2018, company announcement.
[28]LONGi, *Annual Report 2018*, pp. 8, 15, and 26, company report, English translation.

E.2 Outcomes under Continued and Stopped Policy Support

The value of an initial supportive policy depends on whether the government continues to support the industry. Comparing continued and stopped support therefore helps assess whether the benefits of support persist as its future path becomes known. Tables E2 and E3 examine accounting outcomes and downside risk, respectively. Throughout, "stopped" means that no further supportive document appears within the specified window, not that the government explicitly announces termination. Because renewal may respond to an industry's prospects, the comparisons describe differences associated with continuation rather than identify its causal effect. Appendix L separately tests whether observable listed-sector performance and competition predict renewal.

Table E2: **Stopped versus Continued Policy Support: Firm Outcomes**
This table compares firm outcomes for eligible nonfinancial A-share firms whose industries' policy support stops or continues. Each observation begins with a new support spell in year t, defined as support in t but none in $t-1$, during 2010–2019. Continued industries receive support again in $t+1$; stopped industries receive none in $t+1$ or $t+2$; episodes that resume only in $t+2$ are excluded. The first two columns report firm-event means in fiscal year $t+2$. The final column reports the stopped-minus-continued difference. Ratios and growth rates are in decimal units. Parentheses contain industry-clustered standard errors for this difference.

Outcome	Stopped $t+2$ mean	Continued $t+2$ mean	Unadjusted stopped minus continued
Effective tax rate	0.226	0.199	0.027
			(0.021)
Cash growth	0.342	0.420	-0.077
			(0.071)
Short-term-loan growth	0.294	0.435	-0.141
			(0.125)
Long-term-loan growth	0.829	1.062	-0.233
			(0.352)
Operating revenue growth	0.161	0.215	-0.054^{*}
			(0.029)
R&D spending growth	0.368	0.319	0.049
			(0.073)
Operating profit growth	-0.375	0.103	-0.479^{**}
			(0.198)
ROA	0.021	0.032	-0.011
			(0.007)

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table E3: **Stopped versus Continued Policy Support: Downside-Risk Changes**
This table compares downside-risk changes for small firms after stopped versus continued policy support. Each firm-event begins with a supportive announcement in month 0; focal announcements occur through April 2020. Continued industries receive another announcement during months 1–12; stopped industries receive none through month 24; renewals first occurring during months 13–24 are excluded. Each measure is the change from months 0–11 to months 12–23. The first two columns average these changes by group; Difference is stopped minus continued. Returns and frequencies are decimals, volatility measures are annualized, semivariance and LPM1 are multiplied by 10^3, and DUVOL is a log variance ratio. Negative differences indicate worse lower-tail returns and maximum drawdown; positive differences indicate worse outcomes for the other measures. Parentheses contain industry-clustered standard errors.

Family	Outcome	Stopped change	Continued change	Difference
Monthly downside risk	Worst monthly return	-0.013^{*}	0.016^{***}	-0.028^{***}
		(0.007)	(0.003)	(0.008)
	Fraction negative months	0.036	0.000	0.035
		(0.030)	(0.005)	(0.030)
Drawdown risk	Maximum drawdown	-0.017	0.020^{***}	-0.037^{**}
		(0.017)	(0.005)	(0.018)
	Frequency of drawdown $\leq -20\%$	0.013	-0.001	0.015
		(0.014)	(0.003)	(0.016)
Daily tail risk	1st-percentile return	-0.000	0.003^{***}	-0.003
		(0.003)	(0.001)	(0.002)
	5th-percentile return	-0.002	0.005^{***}	-0.007^{***}
		(0.002)	(0.001)	(0.002)
	Expected shortfall, 5%	-0.001	0.004^{***}	-0.005^{***}
		(0.002)	(0.001)	(0.002)
	Frequency of $r \leq -5\%$	0.003	-0.011^{***}	0.014^{***}
		(0.004)	(0.002)	(0.004)
	Frequency of $r \leq -10\%$	0.001	-0.003^{***}	0.004^{***}
		(0.001)	(0.000)	(0.001)
Lower-partial risk	Lower-partial volatility	0.006	-0.024^{***}	0.030^{***}
		(0.011)	(0.005)	(0.010)
	Semivariance below zero ($\times 10^3$)	0.035	-0.081^{***}	0.116^{***}
		(0.030)	(0.015)	(0.029)
	LPM1 below zero ($\times 10^3$)	0.433	-0.687^{***}	1.120^{***}
		(0.321)	(0.169)	(0.324)
Residual downside risk	CAPM downside residual volatility	0.010	0.001	0.009
		(0.007)	(0.003)	(0.006)
Crash-risk diagnostic	DUVOL	0.011	-0.075^{***}	0.085^{**}
		(0.033)	(0.003)	(0.034)

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Appendix F Event Study: Abnormal Returns and Factor Loadings

This appendix evaluates the short-term behavior of asset prices around announcement events. The main takeaway is that supportive-policy announcements do not produce a discrete positive price response, consistent with policy issuance being anticipated or revealing information gradually.

Table F1 reports daily and monthly cumulative abnormal returns for small firms around all supportive-policy announcements. Expected returns are estimated using a market model. The monthly analysis also excludes events whose estimation history overlaps the July 2014–February 2016 Chinese stock-market bubble and crash.

For each firm-event, the cumulative abnormal return is the sum of $AR_{i,\tau} = R_{i,\tau} - \widehat{a}_i - \widehat{b}_i R_{m,\tau}$ over the event window. We average these cumulative returns equally across eligible firm-events. Standard errors are clustered by publication date for daily returns and by policy month for monthly returns.

Table F1: **Short-Horizon Event Studies Around Supportive-Policy Announcements**
This table reports market-model cumulative abnormal returns for small firms around all supportive-policy announcements. Daily windows use trading days; monthly windows use calendar months. The daily estimation window is $[-255, -6]$, with at least 120 observations. The monthly estimation window is $[-41, -6]$, with at least 24 observations. The restricted monthly column excludes events whose nominal $[-41, -6]$ estimation history intersects July 2014–February 2016. Abnormal returns subtract the estimated market-model intercept and market component, are summed over each window, and are averaged equally across firm-events. CARs and clustered standard errors (in parentheses) are in decimals.

	Daily Market Model	*Monthly Market Model*	
Window	Full sample	Full sample	Excluding 2014–2016 estimation-history overlap
(0)	0.000	−0.002	−0.002
	(0.001)	(0.007)	(0.009)
$(0, +1)$	0.000	−0.005	0.007
	(0.002)	(0.011)	(0.017)
$(-1, +1)$	−0.004*	−0.020	−0.001
	(0.002)	(0.014)	(0.021)
$(0, +3)$	0.000	−0.018	−0.009
	(0.003)	(0.017)	(0.029)
$(-3, +3)$	−0.005	−0.057**	−0.040
	(0.004)	(0.023)	(0.029)
$(0, +5)$	0.004	−0.037*	−0.015
	(0.004)	(0.021)	(0.031)
$(-5, +5)$	−0.002	−0.093***	−0.042
	(0.005)	(0.032)	(0.038)

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Table F2 compares CAPM market betas estimated before and after supportive-policy announcements, complementing the downside-risk analysis in Section 4.1. We report all announcements and first or restarted support episodes, using nonoverlapping estimation windows separated from the announcement month. The estimated changes in market beta are statistically insignificant at conventional levels.

Table F2: **Changes in CAPM Market Beta Around Policy Support**
This table compares small firms' CAPM betas estimated from daily excess returns in nonoverlapping calendar-month windows before and after supportive-policy announcements. The 60-month specification compares $[-66, -7]$ with $[6, 65]$ and requires at least 600 daily observations in each period; the 36-month specification compares $[-42, -7]$ with $[6, 41]$ and requires at least 360. First announcements are first or restarted support episodes following more than three years without support. The Pre and Post columns report mean firm-event betas; Δ beta is post minus pre. Standard errors and p-values use industry clustering.

Window	Event sample	Pre beta	Post beta	Δ beta	SE	p-value	Firm-events
60 months	All announcements	1.150	1.172	0.022	(0.019)	0.255	4,787
60 months	First announcements	1.076	1.179	0.104	(0.062)	0.107	261
36 months	All announcements	1.209	1.209	0.000	(0.016)	0.999	8,572
36 months	First announcements	1.153	1.134	−0.018	(0.033)	0.586	340

Appendix G Policy-Portfolio Returns in Adverse States

This appendix provides the implementation details for the adverse-state analysis in Section 6.1, reported in Table 6 and Figure 5. We describe the regression specification, state definitions, and episode dating. These tests examine whether policy-portfolio returns are low in candidate high-marginal-utility states, using conventional financial-market conditions, household and demand conditions, and dated changes in the policy regime.

G.1 Regression Specification

For each state variable Z_t, we estimate a separate regression,

$$R_t^S - R_t^N = \alpha + \boldsymbol{\beta}' F_t^{CH4} + \gamma Z_t + \varepsilon_t, \tag{14}$$

where $R_t^S - R_t^N$ is the value-weighted return spread between supported and non-supported small firms, and F_t^{CH4} contains the market, size, value, and turnover factors.[29] Each Z_t is an adverse-state indicator, so γ is the factor-adjusted adverse-minus-normal monthly return difference. Coefficients and standard errors are in monthly decimal returns. The two PBOC regressions assign each quarterly adverse-state indicator to its three constituent months and cluster standard errors by calendar quarter. The remaining regressions report conventional OLS standard errors. For every state indicator, we use a one-sided test of $H_0 : \gamma \geq 0$ against $H_1 : \gamma < 0$.

G.2 Conventional Financial-Market States

We begin with conventional financial-market state variables: monthly market excess returns, market volatility, and economic policy uncertainty (EPU). These variables capture weak market performance and heightened uncertainty. Monthly market volatility is the within-month standard deviation of daily market excess returns. The EPU measures are the raw Mainland newspaper and *South China Morning Post* indexes, described in Appendix I.

[29] These tests use exact-next-calendar-month returns and renormalize weights over observed returns. The baseline portfolios in Table 5 use next-observed returns and retain formation weights when returns are missing.

We define an adverse-quartile dummy for each variable and report the results in Panel A of Table 6. The market indicator equals one when the monthly market excess return is at or below its 25th percentile. The volatility and EPU indicators equal one when their respective series are at or above the 75th percentile. All other months are coded zero.

G.3 Household and Demand States

We also explore household and demand state variables that capture consumer confidence, manufacturing demand, and households' assessments of income and employment conditions. Consumer confidence is the OECD series distributed by FRED as `CSCICP02CNM460S`. Manufacturing PMI is the official National Bureau of Statistics (NBS) nationwide index. Both series are monthly. The quarterly income- and employment-perception indexes come from the People's Bank of China (PBOC) Urban Depositor Survey. Lower values of these four published indexes indicate worse conditions.

Consistent with Appendix G.2, we define an adverse-quartile dummy for each published index in levels. The dummy equals one when the index is at or below its lower sample quartile, including ties, and zero otherwise. For the PBOC measures, the quarterly dummy is assigned to each of the three months in that quarter. Panel B of Table 6 reports the results.

G.4 Policy-Regime Episodes

In China's discretion-based system, market declines and weak demand may raise expectations of government support. We construct dated policy-regime episodes to examine returns when the direction or continuation of that support is itself uncertain.

For an externally grounded benchmark, we first draw on IMF accounts of five China-wide stress episodes: the June 2013 interbank liquidity squeeze, June 2015 equity and exchange-rate turbulence, June 2018 trade and deleveraging stress, January 2020 COVID shock, and July 2021 property and regulatory tightening. We classify 2015, 2018, and 2021 as policy-regime episodes because these involved changes in the direction, credibility, or continuation of government policy. The IMF accounts describe the mid-2015 equity correction and August exchange-rate-policy change (2015 Article IV), the 2018 transition between deleveraging and renewed growth support (2019 Article IV), and the 2021 property-financing and regulatory

tightening (2021 Article IV). The IMF-motivated indicator equals one in June–August 2015, June–August 2018, and July–September 2021.

To reflect China's institutional context, we also define a China-specific episode set covering uncertainty about political governance and the implementation or durability of policy. This set contains four dated developments:

1. **July–September 2014: anti-corruption escalation.** We date this episode from the July 29 announcement of the investigation of Zhou Yongkang (official announcement).

2. **June–August 2018: trade-war escalation.** We date this episode from the June 15 announcement of the tariff list for Chinese products (dated tariff announcement).

3. **August–October 2020: property-financing rules.** We date this episode from the August 20 meeting associated with the "three red lines" financing framework for property developers (official account of the financing framework).

4. **November 2020–January 2021: platform and technology regulation.** This episode begins with the November 3 suspension of Ant Group's listing and the November 10 platform-antitrust draft (SSE announcement; SAMR consultation notice).

Each episode includes its onset month and the following two months; all other sample months are coded zero. The results are reported in Panel C of Table 6.

Appendix H Additional Tests of Investor Attention and Extrapolative Beliefs

We examine portfolio returns following rapidly publicized policy support and margin financing around policy announcements. These tests address the investor-attention and extrapolative-beliefs explanations discussed in Sections 7.1.1 and 7.1.2, respectively.

H.1 Rapid Broadcast Coverage and Portfolio Returns

We obtain *Xinwen Lianbo* news titles and transcripts from the GovOpenData archive, a third-party archive of CCTV's daily news program. The collected archive spans January 2007–July 2026. We parse each daily page into individual news reports, retaining the broadcast date, report title, transcript, and source link.

We use machine-assisted matching to link these reports to State Council documents. For each document, we search reports dated within 30 calendar days before or after either its publication date or its origination date. We standardize punctuation and spacing and compare document numbers, full policy titles, and distinctive title phrases with report titles and transcripts. Accepted matches contain this document-specific text or a documented paraphrase identifying the policy. Title similarity is used to rank candidate reports. When several reports match a document, we select the earliest accepted report and measure its timing relative to the document's publication date. For accepted reports dated February 3, 2016 onward, we also verify their titles and broadcast dates against CCTV's official daily archive.

The pricing sample contains 275 supportive documents published through November 2021. Rapid coverage means a matched report available by three calendar days after publication, including reports before publication. There are 161 such documents (58.5%): 138 first reported on days 0–3 and 23 reported earlier. Of the remaining 114 documents, 109 have no accepted match and five are first matched after day 3.

We use the small-firm stock screens, weighting schemes, and CH4 factors from the main policy-portfolio analysis (Table 5), with an additional sort by rapid broadcast coverage. At the end of formation month t, a supportive document is active if it was published in months $t-11$ through t. An industry enters the any-rapid portfolio if at least one active document has rapid coverage, including industry-months with both rapid and nonrapid documents. The no-rapid portfolio requires active support but no rapidly covered document. These groups partition supported industry-months. Table H1 reports the two portfolios and their direct return spread.

Table H1: **Rapid *Xinwen Lianbo* Coverage and Small-Firm CH4 Alpha**
This table reports monthly CH4 alphas in decimal returns, with standard errors in parentheses. Any-rapid includes supported industry-months with at least one active document matched to a broadcast available by day 3; no-rapid requires active support without a rapid match. Active documents are those published in months $t-11$ through t.

	Equal-weighted	Value-weighted
Any rapid coverage	0.006***	0.006***
	(0.002)	(0.002)
No rapid coverage	0.001	0.001
	(0.005)	(0.005)
Any rapid − no rapid	0.006	0.006
	(0.005)	(0.005)

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

H.2 Margin Financing Around Policy Announcements

This subsection tests whether margin-financed trading increases around the issuance of supportive industrial policies and subsequently unwinds, as suggested by the extrapolative-beliefs explanation in Section 7.1.2. We use CSMAR's daily stock-level margin-financing data from March 31, 2010–December 31, 2021. The sample includes stocks of all sizes satisfying the prior-month listing-age, trading-history, and capitalization screens.

Each panel of Figure H1 plots changes in margin-financed trading from quarter -4 through quarter $+4$ around a policy announcement, using event-time coefficients and 95% confidence intervals from the following panel regression:

$$Financing_{e,t,q} = \sum_{\substack{k=-4 \\ k\neq -1}}^{+4} \beta_k D^k_{e,t,q} + a_e + a_t + \epsilon_{e,t,q},$$

where e indexes industry-announcement events, t denotes the calendar day, and q is the quarter relative to the announcement. The indicators $D^k_{e,t,q}$ identify these relative quarters, with quarter -1 omitted as the reference period. The terms a_e and a_t are event and calendar-day fixed effects.

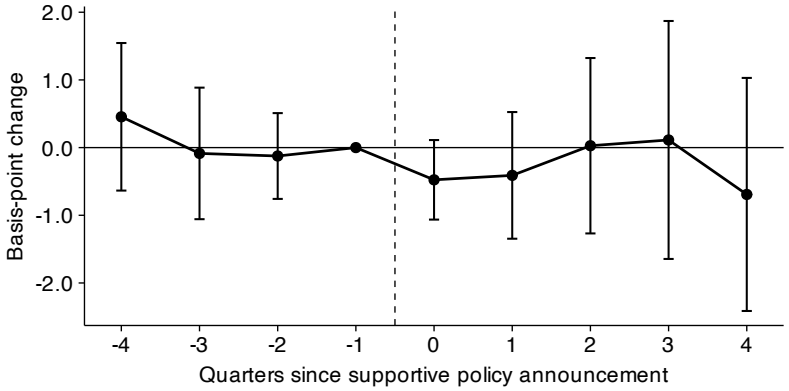
In Panel A, financing is measured by gross margin-financed purchases, defined as total purchases in the industry divided by the same stocks' total lagged float market capitalization. In Panel B, financing is measured by net flow, defined as total net purchases (purchases minus repayments) divided by the same denominator. Both measures are multiplied by 10,000 to express them in basis points. Capitalization comes from each stock's previous observed trading record. We winsorize at the 1st and 99th percentiles across industries on dates with at least 20 finite observations.

We date each supportive policy announcement to the first trading day on or after publication. The nine quarters from -4 through $+4$ use the following calendar-day bins: $[-365, -274]$, $[-273, -183]$, $[-182, -92]$, $[-91, -1]$, $[0, 91]$, $[92, 182]$, $[183, 273]$, $[274, 364]$, and $[365, 456]$. The full window must fall within the data period. Each quarter must contain at least 40 exchange-open days, with valid gross and net financing observations on at least 80% of those days. These requirements leave 485 events in 43 industries, dated April 27, 2011–September 28, 2020. Standard errors are clustered by industry and calendar day.

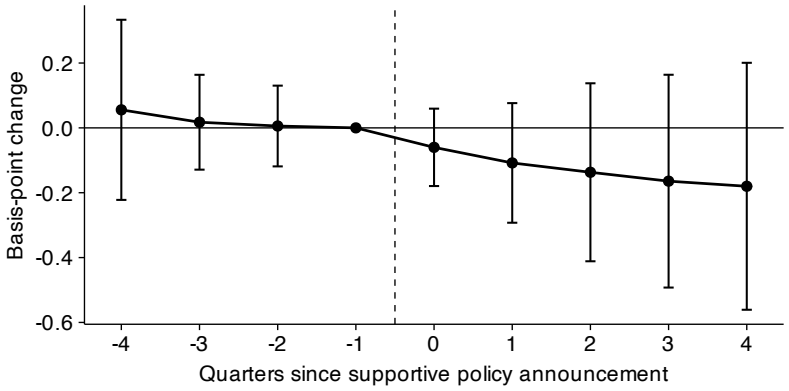
Neither gross margin-financed purchases nor net financing flow increases unusually around supportive-policy announcements.

Figure H1: **Margin Financing Around Supportive-Policy Announcements**
This figure plots changes in margin-financed trading from four quarters before to four quarters after supportive-policy announcements. Panel A plots gross margin-financed purchases; Panel B plots signed net financing flow, defined as purchases minus repayments. Each industry-day outcome equals financing activity summed over all observed stocks, divided by their lagged float market capitalization and expressed in basis points. Outcomes are winsorized by date at the 1st and 99th percentiles. Regressions include event and calendar-date fixed effects. The points are event-time coefficients relative to quarter −1; 95% confidence intervals are based on standard errors clustered by industry and calendar date. The calendar-day bins and coverage rule are defined in Appendix H.2.

Panel A: Gross Margin-Financed Purchases



Panel B: Net Financing Flow



Appendix I Construction of Additional Risk Factors

This appendix describes the additional factors used in Section 7.3. We construct factor-mimicking portfolios for innovations in two China-specific economic policy uncertainty (EPU) indices and an R&D investment factor. The first EPU index, developed by Baker et al. (2016), uses articles from the *South China Morning Post* (SCMP). The second, developed by Davis et al. (2019), uses articles from *Renmin Daily* and *Guangming Daily* (RDGD).

Following Brogaard and Detzel (2015), we extract innovations from each EPU index by estimating

$$EPU_t = a + b_1 EPU_{t-1} + b_2 EPU_{t-2} + b_3 EPU_{t-3} + c_1 MKT_{t-1} + \epsilon_t^{\text{EPU}}. \tag{15}$$

Here, MKT_{t-1} is the lagged monthly Chinese market excess return. We project the estimated innovations onto 25 Chinese size and momentum portfolio returns, denoted by R_t^e:

$$\hat{\epsilon}_t^{\text{EPU}} = a_{EPU} + b'_{EPU} R_t^e + \eta_t^{EPU}. \tag{16}$$

The corresponding factor-mimicking portfolio is

$$EPU_F = \hat{b}'_{EPU} R_t^e. \tag{17}$$

The first two rows of Panel B in Table 8 report CH4 regressions augmented with each EPU factor.

Following Leung et al. (2020), we construct the R&D factor from six value-weighted portfolios sorted on firm size and R&D intensity, measured as R&D expenditure divided by market value. Firms with positive R&D spending are sorted annually into two size groups and three R&D-intensity groups. The factor holds equal-weighted positions in the small-high and large-high portfolios and shorts equal-weighted positions in the small-low and large-low portfolios. The final row of Panel B in Table 8 adds this factor to the CH4 model.

Appendix J Additional Figures and Tables

Table J1: **Policy Portfolio Returns for Small Firms: Full Factor Loadings**
This table reports full asset-pricing regressions for equal- and value-weighted portfolios of supported small firms. Small firms are below the median market capitalization. A firm enters the supported portfolio when the corresponding industry received a supportive State Council policy document during the preceding 12 months. The portfolio is rebalanced monthly. Returns are evaluated using the CAPM, Carhart, and CH4 models from June 2009 to December 2021. Standard errors are reported in parentheses.

	Dependent variable: Monthly portfolio returns					
	Equal-weighted			Value-weighted		
	CAPM	Carhart	CH4	CAPM	Carhart	CH4
const (alpha)	0.012**	0.006***	0.006***	0.010**	0.005***	0.006***
	(0.005)	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)
mktrf	1.142***	0.977***	0.996***	1.136***	0.972***	0.988***
	(0.069)	(0.018)	(0.021)	(0.068)	(0.018)	(0.021)
smb		0.992***	1.084***		0.966***	1.081***
		(0.032)	(0.036)		(0.032)	(0.036)
hml/vmg		−0.171***	−0.319***		−0.198***	−0.318***
		(0.047)	(0.043)		(0.046)	(0.043)
mom		−0.130***			−0.114***	
		(0.023)			(0.023)	
pmo			0.069*			0.064*
			(0.037)			(0.037)
R^2	0.650	0.977	0.975	0.654	0.978	0.975
Adjusted R^2	0.648	0.976	0.974	0.652	0.977	0.974

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Table J2: **Policy Portfolio Returns for Large Firms**
This table reports the returns to a policy portfolio restricted to large firms, defined as those above the median market capitalization in each month. The portfolio is constructed and evaluated using the same methodology as in the main analysis. Estimated alphas are generally smaller than the corresponding estimates for small firms. The equal-weighted CH4 alpha is positive and statistically significant, but the result is not robust to value weighting. Standard errors are reported in parentheses.

	Dependent variable: Monthly portfolio returns					
	Equal-weighted			Value-weighted		
	CAPM	Carhart	CH4	CAPM	Carhart	CH4
const (alpha)	0.003	0.000	0.005**	0.001	0.002	0.003
	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
mktrf	1.077***	0.972***	0.971***	0.991***	0.978***	0.968***
	(0.045)	(0.025)	(0.027)	(0.030)	(0.030)	(0.032)
smb		0.425***	0.545***		−0.112**	−0.035
		(0.044)	(0.047)		(0.053)	(0.055)
hml/vmg		−0.353***	−0.347***		−0.217***	−0.115*
		(0.065)	(0.056)		(0.077)	(0.065)
mom		−0.023			−0.059	
		(0.032)			(0.038)	
pmo			−0.039			−0.102*
			(0.048)			(0.056)
R^2	0.794	0.940	0.943	0.882	0.888	0.898
Adjusted R^2	0.792	0.939	0.941	0.881	0.885	0.895

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Table J3: **Controlling for the Small-Firm Anomaly: Comparing Supported to Non-Supported Firms**

This table evaluates whether the observed abnormal returns to supported firms simply reflect the small-firm anomaly. We construct a long-short portfolio that holds small firms receiving policy support and shorts small firms without support. We then estimate alphas under standard asset pricing models. The difference in returns is positive across all specifications and is strongest under the CH4 model, which explicitly controls for the small-firm effect. This suggests that our results are not driven by generic small-firm premia. For brevity, only the coefficients of interest are reported. Standard errors are reported in parentheses.

	Dependent variable: Monthly policy portfolio returns					
	Equal-weighted			Value-weighted		
	CAPM	Carhart	CH4	CAPM	Carhart	CH4
const (alpha)	0.003*	0.002*	0.006***	0.003*	0.002*	0.006***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)
R^2	0.032	0.354	0.324	0.021	0.368	0.320
Adjusted R^2	0.026	0.337	0.306	0.015	0.351	0.301

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table J4: **Policy Abnormal Returns Are Not Driven by Investor Inattention**
This table tests whether the abnormal returns to policy-supported firms are driven by investor inattention. Each month, we sort small supported stocks based on proxies for industry-level attention and form a long-short portfolio that takes long positions in low-attention industries and short positions in high-attention industries. If inattention explains the observed returns, this portfolio should earn positive alphas. Panel A uses median news intensity, defined as the share of daily news articles related to the industry, while Panel B uses median trading volume share. We find no support for the inattention hypothesis: estimated alphas are statistically insignificant or negative. The news analysis covers March 2010–July 2016 and is limited by data availability in CSMAR. For brevity, only the constant (alpha) term is reported. Standard errors are reported in parentheses.

Panel A: Attention portfolios sorted on news intensity						
	Dependent variable: Monthly attention portfolio return					
	Equal-weighted			Value-weighted		
	CAPM	Carhart	CH4	CAPM	Carhart	CH4
const (alpha)	0.002	0.001	−0.001	0.002	0.001	0.000
	(0.002)	(0.002)	(0.004)	(0.002)	(0.002)	(0.004)
R^2	0.009	0.055	0.018	0.017	0.068	0.023
Adjusted R^2	−0.005	0.002	−0.037	0.003	0.016	−0.032
Panel B: Attention portfolios sorted on trading volume						
	Dependent variable: Monthly attention portfolio return					
	Equal-weighted			Value-weighted		
	CAPM	Carhart	CH4	CAPM	Carhart	CH4
const (alpha)	−0.004*	−0.003	−0.010***	−0.005**	−0.004*	−0.009***
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)
R^2	0.004	0.031	0.037	0.001	0.024	0.032
Adjusted R^2	−0.003	0.004	0.011	−0.006	−0.003	0.005

*$p < 0.10$, **$p < 0.05$, and ***$p < 0.01$.

Table J5: **Policy Portfolio Returns by Firm-Level R&D Growth**
This table examines whether the abnormal returns of supported firms are correlated with firm-level R&D growth. Each month, we sort firms in the policy portfolio into high and low R&D growth groups, based on whether their R&D spending growth over the past year is above or below the median. We then estimate monthly alphas for each group using the CAPM, Carhart, and CH4 models. This analysis motivates the construction of a formal R&D risk factor in the main robustness tests. For brevity, only the coefficients of interest are reported. Standard errors are reported in parentheses.

	Dependent variable: Monthly policy portfolio returns					
	Low R&D Growth			High R&D Growth		
	CAPM	Carhart	CH4	CAPM	Carhart	CH4
const (alpha)	0.009*	0.003	0.004*	0.016***	0.009***	0.015***
	(0.005)	(0.002)	(0.002)	(0.005)	(0.002)	(0.003)
R^2	0.605	0.950	0.936	0.597	0.943	0.883
Adjusted R^2	0.602	0.948	0.934	0.594	0.941	0.880

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table J6: **Supportive Policy and Industry IPO Intensity**
This table examines the relationship between supportive policy and subsequent IPO intensity at the industry-year level. Columns (1)–(2) define IPO intensity as the industry's share of all sample listings in year t. Columns (3)–(4) define it as new listings in year t divided by the prior listed-firm count, measured as the number of sample firms in the industry listed by the end of $t-1$. Both outcomes are ratios winsorized at the 1st and 99th percentiles of the 2009–2021 panel. The policy indicator equals one if the industry's first observed supportive policy occurred in $t-2$ or $t-3$, and zero otherwise. First observed support is the earliest supportive document published in our sample. Columns (2) and (4) exclude 2013, which was affected by a regulatory pause in IPOs. All regressions include industry and year fixed effects, log(1 + prior listed-firm count), the IPO count in $t-4$, and log(1 + aggregate listed-firm market capitalization in $t-4$). The two $t-4$ controls precede either first-support timing. Industry-clustered standard errors are in parentheses.

	IPO share		IPO rate	
	(1) Full sample	(2) Excl. 2013	(3) Full sample	(4) Excl. 2013
First support in $t-2$ or $t-3$	0.005***	0.002**	0.041***	0.035**
	(0.002)	(0.001)	(0.013)	(0.015)
Industry fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	810	729	797	718

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Appendix K Implied Cost of Capital Tests

This appendix examines whether policy support is associated with changes in forward-looking measures of the cost of equity. The tests use two implementations of the Gebhardt–Lee–Swaminathan (GLS) residual-income model (Gebhardt et al., 2001). The first uses analyst earnings forecasts, while the second replaces analyst forecasts with earnings predicted from accounting information. The analysis asks whether ICC changes following policy support and whether this response differs between small and large firms. We find no evidence that ICC increases following fresh support. When support continues into a second consecutive year, however, small-firm ICC declines, consistent with financing benefits becoming more apparent as support continues.

K.1 ICC construction

Let $P_{i,t}$ denote the June stock price and $B_{i,t}$ book value per share. For each firm-year, ICC is the discount rate $r_{i,t}$ that equates price to current book value plus discounted expected residual income:

$$P_{i,t} = B_{i,t} + \sum_{\tau=1}^{12} \frac{(FROE_{i,t+\tau} - r_{i,t})B_{i,t+\tau-1}}{(1 + r_{i,t})^{\tau}}. \tag{18}$$

Forecast return on equity is constructed from explicit earnings forecasts for the first three fiscal years. It then fades linearly to the historical industry target through year 12. Book value evolves using forecast earnings and a payout ratio based on dividends and earnings. Valid ICC roots lie in $(0, 1]$, and each ICC measure is winsorized at the 1st and 99th percentiles.

The analyst measure uses median EPS forecasts for fiscal years t, $t+1$, and $t+2$ available by June 30. Following Hou et al. (2012), the model-based measure predicts total RMB earnings for the same three fiscal years in rolling ten-year windows using only information observable by the June measurement date, then converts predicted earnings to EPS. Both measures use June prices, fiscal-year $t-1$ accounting information, and the same payout and terminal-value assumptions.

The common sample contains 22,361 firm-years. Both measures are winsorized again at the 1st and 99th percentiles within this common sample. The analyst and model-based measures have means of 0.056 and 0.042 and standard deviations of 0.028 and 0.025, respectively.

Their Pearson and Spearman correlations are 0.879 and 0.875.

K.2 Policy support and changes in ICC

The preferred annual specification maps supportive policies issued during calendar year t to the change in ICC from June of year t to June of year $t+1$. This midpoint alignment approximates the average 12-month post-announcement window in the monthly portfolio analysis, but it does not assign a separate ICC observation to each policy date. The support indicator equals one if the firm's industry receives at least one supportive policy during year t. We also use $SupportStr$, defined as the number of supportive policy–industry links for the industry divided by the number of unique supportive documents issued during the year.

For ICC changes, we estimate

$$\Delta ICC_{i,t+1} = \beta_1 Support_{j,t} + \beta_2 Small_{i,t} + \beta_3 Support_{j,t} \times Small_{i,t} + \Gamma' X_{i,t} + \alpha_j + \delta_t + \epsilon_{i,t+1}, \quad (19)$$

where j denotes the firm's industry. The controls are log market capitalization, book-to-market, beta, 12-month momentum, ROE, leverage, state ownership, and listing age. Analyst-ICC regressions also control for the number of FY1 analysts. All specifications include industry and year fixed effects, and standard errors are clustered by industry. In the interaction specifications, β_1 is the large-firm effect and β_3 is the difference between the small- and large-firm effects.

Table K1 reports the change regressions. Using the support indicator, large-firm ICC declines by 32.2 basis points with analyst ICC and 21.8 basis points with model-based ICC. The small-minus-large interactions are 18.6 and 13.4 basis points, respectively. The interactions based on support strength are also positive and significant at the 1% level. These results show that small firms receive less of the ICC decline associated with support among large firms.

Table K1: **Policy Support and Implied Cost of Capital**
This table reports separate firm-level regressions of the June-to-June change in ICC on a support indicator (Panel A) or *SupportStr* (Panel B), each interacted with *Small*. The indicator equals one if the firm's industry receives any supportive policy during calendar year t. *SupportStr* is the number of supportive policy–industry links divided by the number of unique supportive documents issued during the year. *Small* identifies firms below the annual median June market capitalization. All regressions include industry and year fixed effects and controls for log market capitalization, book-to-market, beta, 12-month momentum, ROE, leverage, state ownership, and listing age; analyst regressions also control for the FY1 analyst count. Standard errors, in parentheses, are clustered by industry.

	Dependent variable: $\Delta ICC_{i,t+1}$	
	Analyst GLS median	Model-based GLS
Panel A: Support indicator		
Support indicator	-0.003***	-0.002**
	(0.001)	(0.001)
Support indicator × Small	0.002**	0.001**
	(0.001)	(0.001)
Panel B: Support strength		
SupportStr	-0.023**	-0.014**
	(0.011)	(0.006)
SupportStr × *Small*	0.019***	0.013***
	(0.005)	(0.004)
Observations	14,455	31,197
Controls	Yes	Yes
Year fixed effects	Yes	Yes
Industry fixed effects	Yes	Yes
Industry-clustered SE	Yes	Yes

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

K.3 Initial and renewed support

We also test whether the ICC response changes when support is repeated. Fresh support denotes support in year t but not $t-1$. The second consecutive year denotes support in both t and $t-1$ but not $t-2$. The regressions also include a separate indicator for support lasting at least three consecutive years. The omitted category is no support in year t.

Table K2 shows no significant small-firm ICC response to fresh support. In the second consecutive support year, the implied small-firm effect is -43.1 basis points using analyst ICC and -44.8 basis points using model-based ICC. In the small-firm-only regressions, model-based ICC falls by 23.0 basis points in the second year, and the decline relative to fresh support is 31.3 basis points. The corresponding analyst estimates are negative but imprecise.

Effects after three or more consecutive years are also imprecise and are not reported in the compact table. This dynamic evidence is consistent with financing benefits becoming more apparent when support is renewed.

Table K2: **ICC Response to Initial and Renewed Policy Support**
This table compares changes in ICC from June of year t to June of year $t+1$ following fresh support and a second consecutive year of support. Fresh support occurs in year t but not $t-1$, while second-year support occurs in both t and $t-1$ but not $t-2$. The omitted state is no support in year t; each regression also includes an unreported indicator for three or more consecutive years of support. Panel A reports the implied small-firm effects from full-sample interaction regressions, while Panel B restricts the sample to small firms. The difference and Wald p-value rows compare second-year with fresh support. All regressions include industry and year fixed effects and controls for log market capitalization, book-to-market, beta, 12-month momentum, ROE, leverage, state ownership, and listing age; analyst regressions also control for the FY1 analyst count. Coefficients and standard errors are expressed in basis points, with standard errors clustered by industry.

	Analyst GLS median	Model-based GLS
Panel A: Full-sample interaction, implied small-firm effect		
Fresh support	-11.2	1.2
	(11.6)	(7.7)
Second consecutive year	-43.1**	-44.8***
	(19.9)	(11.3)
Second year minus fresh	-31.9	-46.0***
Wald p-value	0.157	< 0.001
Panel B: Small-firm-only regression		
Fresh support	-3.9	8.3
	(10.7)	(7.3)
Second consecutive year	-16.1	-23.0**
	(22.3)	(10.9)
Second year minus fresh	-12.2	-31.3***
Wald p-value	0.607	0.009
Full-sample observations	14,455	31,197
Small-firm-only observations	3,599	15,096
Controls	Yes	Yes
Industry and year fixed effects	Yes	Yes
Industry-clustered SE	Yes	Yes

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Appendix L Observable Industry Information and Policy Renewal

The divergence between short and longer support spells may reflect industry triage: the government could initially support several candidate industries and then renew support only for those that reveal stronger prospects. A related possibility is that authorities are more likely to renew support in industries with stronger product-market competition. We therefore examine observable performance and predetermined competition as predictors of renewal. This is a selection test and is distinct from asking whether the effects of support are stronger in more competitive industries.

To match this conjecture, the renewal risk set is restricted to the first or second consecutive year of a support spell during 2010–2020. The unrestricted early-stage risk set contains 140 industry-years, including 68 renewals and 72 stops across 49 industries. Requiring complete listed-sector performance data and coverage of at least five firms leaves 90 industry-years, including 42 renewals and 48 stops across 36 industries. The competition measures are available for 85 of these observations, including 40 renewals and 45 stops across 33 industries. The performance screen retains 61.8% of renewals and 66.7% of stops from the unrestricted risk set; the difference in retention is statistically insignificant (Fisher's exact $p = 0.599$).

Panel A of Table L1 reports unadjusted comparisons. None of the five renewed-minus-stopped differences is statistically significant. Panel B conditions on calendar year, support age, and observed prior support history. The three performance measures are jointly insignificant in every specification. A one-standard-deviation increase in PCM competition is associated with an 8.00-percentage-point increase in renewal probability, and the corresponding estimate for HHI competition is 4.91 percentage points, but neither estimate is statistically significant. Thus, the point estimates are consistent with competition-based selection, but the evidence is imprecise. The results provide no robust evidence that observable listed-sector performance or competition explains early policy renewal. They do not rule out triage based on government information, nonlisted firms, employment, technological capacity, or strategic considerations that are not observed in the listed-firm data.

Table L1: **Industry Performance, Competition, and Early Policy Renewal**

This table tests whether observable listed-sector performance and competition predict policy renewal during the first two years of a support spell. The unit is a two-digit CSRC industry-year with support in year t; renewed equals one if support continues in $t+1$. Panel A reports unadjusted means and renewed-minus-stopped differences. Panel B reports linear probability models. Column (1) includes the three performance measures. Columns (2) and (3) separately add PCM and HHI competition; performance coefficients are included but omitted from these columns. All models include calendar-year fixed effects and indicators for the second support year and an observed recurrent support spell. Performance is measured in fiscal year $t-1$; sales and margin changes require five listed firms in $t-2$ and $t-1$, and returns require five return-bearing firms in $t-1$. The competition measures average fiscal years $t-2$ through $t-4$, require at least two usable five-firm years, and are oriented so higher values denote greater competition. Continuous measures are winsorized at the 1st and 99th percentiles and standardized within each regression sample. Parentheses contain industry-clustered standard errors.

Panel A: Unadjusted characteristics by subsequent renewal status

	Renewed	Stopped	Renewed − stopped
Performance			
Sales growth (%)	25.16	23.21	1.95
			(7.46)
Operating-margin change (pp)	0.16	0.48	−0.32
			(0.93)
Latest industry return (%)	10.57	7.18	3.38
			(9.26)
Predetermined competition			
PCM competition (pp)	−5.22	−7.56	2.34
			(1.89)
HHI competition (pp)	86.85	82.20	4.65
			(3.33)

Panel B: Predicting early policy renewal

	(1) Performance	(2) + PCM	(3) + HHI
Sales growth	−8.17*		
	(4.49)		
Operating-margin change	−0.55		
	(4.04)		
Latest industry return	−5.57		
	(9.60)		
Predetermined competition			
PCM competition		8.00	
		(5.13)	
HHI competition			4.91
			(4.42)
Joint p-value, performance measures	0.325	0.559	0.251
Observations	90	85	85
Renewals / stops	42 / 48	40 / 45	40 / 45
Industries	36	33	33
Calendar-year fixed effects	Yes	Yes	Yes
Support-age/history controls	Yes	Yes	Yes

$^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Appendix M Adverse Policy Realizations in Supported Industries

This appendix illustrates two adverse policy realizations in previously supported industries. For each case, we first describe the policy support in place before the shock and then the adverse policy action. Table M1 reports the corresponding industry returns. These cases complement the systematic portfolio evidence in Section 6.1.

M.1 Internet and Software Services, November 2020

On July 30, 2020, a State Council implementation opinion described industrial-internet and e-commerce platforms as infrastructure for entrepreneurship, promoted platform-enabled online business, and called for support for internet-platform enterprises.[30] However, the policy direction changed on November 10, when the State Administration for Market Regulation released a consultation draft applying antitrust rules to the platform economy. The draft addressed platform agreements, abuse of dominance, restricted dealing, and merger review.[31]

Table M1 reports the November 10–11 returns for the Internet and software industries separately. Both industries declined and underperformed the market.

M.2 Education, July 2021

In July 2021, the education industry (CSRC 82) was covered by policy support under the May 14 implementing regulation for the Private Education Promotion Law. The regulation provided subsidies, tax and land preferences, special funds, and credit and financial products for private education.[32] Double Reduction directive was then released on July 24 after its restrictions circulated publicly on July 23. It required existing compulsory-education subject-tutoring providers to become nonprofit organizations, prohibited their access to listed

[30] See the July 30, 2020 implementation opinion. A January 17, 2020 State Council document on national new areas also supported these industries; see the official guidance.

[31] See the SAMR consultation notice. The Shanghai Stock Exchange had suspended Ant Group's listing on November 3 following regulatory interviews and a material change in the financial-technology regulatory environment; see the Shanghai Stock Exchange announcement.

[32] See the May 14, 2021 regulation and the Ministry of Education explanation.

financing and capitalization, and barred tutoring on weekends and school holidays.[33]

Table M1 reports the education industry's July 23–27 return. Together, the cases illustrate adverse policy changes in previously supported industries.

Table M1: **Industry Returns Around Adverse Policy Realizations**
This table reports raw stock returns for the previously supported industries discussed in Appendix M. Panel A covers November 10–11, 2020, beginning with the release of the platform-antitrust consultation draft. Panel B covers July 23–27, 2021, from the public circulation of the Double Reduction restrictions through the first trading days after the formal directive. Equal-weighted returns average constituent buy-and-hold returns; value-weighted returns use prior-month market capitalization. The market return is the CH4 market total return. Returns are in decimal units.

Industry	Event window	EW return	VW return	Market return
Panel A: Platform-antitrust consultation draft				
Internet and related services	Nov. 10–11, 2020	−0.036	−0.036	−0.019
Software and IT services	Nov. 10–11, 2020	−0.040	−0.041	−0.019
Panel B: Double Reduction directive				
Education	July 23–27, 2021	−0.198	−0.195	−0.061

[33]See the July 23 contemporaneous report and the July 24 official directive.