Working Paper Series

SSRN Draft – Central-Bank Research

Satellite Nowcasting and the Effectiveness of Monetary Policy

Yue Dai

Institute of Finance Research, People's Bank of China

[dyue@pbc.mail.cn]

June 2026

Please cite this paper as:

Dai, Yue (2026). "Satellite Nowcasting and the Effectiveness of Monetary Policy," SSRN Working Paper.

Abstract

This paper studies whether satellite imagery can improve monetary-policy transmission by supplying timely public information about the state of the economy. I construct a satellite-based nowcasting system that maps nighttime lights, land surface temperature, and industrial-emission measures to firm-level production signals and aggregate macroeconomic indicators. The remote-sensing panel covers more than 1,000 manufacturing firms; based on the 1,346 satellite-observable firms in the combined model and the 2015–2023 monthly sample, the processing pipeline implies roughly 145,000 firm-month image chips and about 872,000 firm-product-month signal cells before lag and rolling-feature expansion. The combined satellite signal predicts firm revenue growth with an accuracy of 81.9 percent and an F1 score of 90.3 percent; at the macro level, satellite-observable firms are highly correlated with GDP and manufacturing PMI, and a satellite classifier predicts whether PMI is above the expansion threshold with 71.3 percent accuracy, rising to 80.0 percent when cloud cover is low. I then develop a Bayesian model in which the policy rate is both an instrument and a signal of central-bank information. More precise nowcasts reduce the private sector's need to infer fundamentals from monetary-policy surprises, thereby attenuating the information-effect component of policy announcements and strengthening the direct interest-rate channel. Empirically, I use cloud-cover-induced variation in satellite observability to measure the validity of big-data nowcasts and estimate state-dependent local projections with high-frequency monetary-policy instruments. The results indicate that better satellite observability weakens the equity-market information effect of policy surprises and increases pass-through to short- and long-term bond yields. The findings suggest that real-time alternative data can function as monetary-policy infrastructure by improving the common information set on which policy expectations are formed.

JEL classification: C53, E52, E58, G14, O33.

Keywords: satellite imagery; nowcasting; monetary-policy transmission; central-bank information effect; local projections; alternative data.

1 Introduction

Central banks make policy under imperfect and unevenly distributed information. The problem is especially acute when official data are published with substantial delay. In China, manufacturing PMI is released only after the reference month closes, most monthly activity indicators are released later, and GDP is available only after the quarter ends. Policy decisions, however, must be forward-looking. The resulting timing mismatch creates an inter-release period in which market participants infer the state of the economy not only from data releases, but also from the actions of the central bank itself.

This paper examines whether satellite imagery can reduce that informational friction. Satellite data have three properties that are valuable for monetary analysis. They are high frequency, spatially granular, and based on physical traces of activity rather than survey responses or accounting reports. Nighttime lights proxy for the intensity of local activity, land surface temperature captures heat generated by production, and atmospheric emissions reflect industrial operation. These signals are available before official macroeconomic statistics and can therefore be used to nowcast the current state of the economy.

The paper makes two arguments. The first is measurement oriented: satellite imagery contains information about current economic activity that is useful at both the firm and aggregate levels. The second is monetary: the value of nowcasting is not limited to internal central-bank forecasting. By improving the public information environment, nowcasting can change how private agents interpret policy surprises. When the private sector has a more precise independent signal of current fundamentals, it relies less on the policy rate as a signal of the central bank's private information. The information-effect component of monetary-policy announcements is therefore reduced.

The empirical analysis proceeds in three steps. First, I build a micro-to-macro satellite nowcasting system from a large panel of firm-level image chips. At the firm level, rolling regularized models use satellite features to predict the direction of revenue growth. At the aggregate level, satellite-observable firms are used to construct indicators for PMI and industrial activity, including mixed-frequency MIDAS specifications. Second, I formalize the information channel in a Bayesian model of monetary-policy transmission. Third, I test the model using local projections with high-frequency monetary-policy instruments. Cloud cover provides plausibly exogenous variation in the information content of satellite imagery: when cloud cover is high, satellite observability is poor and nowcasts are less informative; when cloud cover is low, satellite observability is better and nowcasts are more precise.

The results are consistent with the proposed mechanism. The combined satellite signal predicts firm revenue-growth direction with 81.9 percent accuracy and an F1 score of 90.3

percent. Satellite-observable firms are strongly representative of aggregate activity: their signal correlates above 0.90 with manufacturing PMI and above 0.94 with GDP. The PMI expansion classifier has 71.3 percent accuracy overall and 80.0 percent accuracy under low-cloud conditions. In the monetary-policy analysis, better nowcast validity weakens the positive equity-market response associated with central-bank information effects and strengthens pass-through to short- and long-term bond yields.

The contribution is to connect two literatures that have largely evolved separately. The nowcasting literature emphasizes the real-time informational content of high-frequency data [Giannone et al., 2008, Bok et al., 2018]. The monetary-policy literature emphasizes that policy surprises may contain information about central-bank beliefs rather than only exogenous shifts in the policy stance [Nakamura and Steinsson, 2018, Jarocinski and Karadi, 2020]. This paper shows that a better nowcasting technology can alter the information structure that gives rise to those effects.

2 Related Literature and Institutional Background

The paper contributes to three strands of research. The first studies real-time macroeconomic measurement. Since Giannone et al. [2008], nowcasting has become a standard approach for extracting the current state of the economy from data that arrive at different frequencies and with different publication lags. Central banks have used dynamic-factor models, bridge equations, and MIDAS regressions to summarize incoming information in a way that is useful for policy deliberations. The contribution here is to move beyond conventional financial and survey indicators by evaluating satellite imagery as a source of real-time information. Satellite data are not merely another high-frequency series. They are generated independently of the statistical reporting system and measure physical activity directly. This makes them especially attractive in periods when survey data are delayed, noisy, or subject to unusual reporting conditions.

The second strand uses alternative data to measure economic activity. Nighttime lights have been used to proxy for local growth, urbanization, and regional development [Henderson et al., 2012, Donaldson and Storeygard, 2016]. More recent work uses remote-sensing data, mobility data, payment flows, shipping records, and online prices to measure activity at higher frequencies than official statistics. The measurement problem in this literature is twofold. First, an alternative signal must have predictive content for real activity. Second, the signal must have a transparent failure mode, so that users know when it is more or less reliable. The satellite setting is useful because cloud cover provides a direct and observable measure of data quality. A signal based on optical or thermal imagery should be less informative when cloud cover is high. This feature allows the paper to separate the information content of the nowcast from general changes in macroeconomic conditions.

The third strand studies central-bank information effects. High-frequency identification treats narrow-window movements in asset prices around policy announcements as monetary-policy surprises, but a policy surprise can contain more than one component. It can reflect an exogenous shift in the expected policy path, and it can reveal the central bank's private information about the economy [Nakamura and Steinsson, 2018, Jarocinski and Karadi, 2020]. The latter component is central for interpreting asset-price responses. A tightening that would ordinarily lower equity prices may instead raise them if investors infer that the central bank has received favorable news about growth. This paper asks whether better public nowcasting reduces the scope for such inference. If the private sector already has a precise signal of current fundamentals, policy announcements should carry less incremental information about the state of the economy.

The institutional setting strengthens the relevance of the question. In a price-based monetary-policy framework, short-term money-market rates play an increasingly important role in transmitting policy to broader financial conditions. At the same time, the official macroeconomic data calendar creates predictable intervals during which the most recent activity data are not yet available. The central bank must act on forward-looking assessments, while private agents observe the action and update beliefs about both policy and fundamentals. In such a setting, the same interest-rate surprise can be interpreted through two lenses: as a change in the stance of policy, and as a signal about the central bank's information set. A real-time nowcast is valuable precisely because it can discipline the second interpretation.

The paper therefore treats nowcasting as part of the monetary-policy transmission environment. This perspective differs from a purely forecasting exercise. A forecast is useful if it predicts the target variable. A policy-relevant nowcast is useful if it changes the mapping from policy announcements to private expectations. The empirical sections below first establish that satellite imagery has nowcasting content, then test whether variation in that content changes the transmission of monetary-policy shocks.

3 A Model of Nowcasting and Policy Information Effects

This section develops the theoretical mechanism before turning to the measurement exercise. The ordering is deliberate. The object of interest is not satellite imagery per se, but whether the availability of a timely and precise nowcast improves the transmission of monetary policy.

The model clarifies why nowcast precision is a state variable for policy effectiveness.

The model is intentionally parsimonious. It abstracts from the full dynamic structure of a New Keynesian economy and focuses on the informational margin that is most relevant for the empirical design. The policy rate can move asset prices through a direct intertemporal-substitution channel, but it can also move asset prices because it reveals information about the central bank's assessment of current fundamentals. The introduction of a public nowcast changes the relative strength of these two channels. The main comparative static is therefore not whether better data make the central bank more informed, but whether better public data reduce the private sector's dependence on the policy action as a signal.

3.1 Environment

Let θ denote the current economic fundamental. The fundamental is not directly observed and has prior distribution

$$\theta \sim N(0, \tau_0^{-1}), \tag{1}$$

where τ_0 is prior precision. Output is determined by the fundamental and the policy rate i:

$$y = \theta - \phi i, \qquad \phi > 0. \tag{2}$$

The central bank sets the policy rate according to a forward-looking rule,

$$i = \psi E^{CB}[\theta \mid \mathcal{I}^{CB}] + \varepsilon_i, \tag{3}$$

where ψ is publicly known and $\varepsilon_i \sim N(0, \tau_i^{-1})$ is a conventional policy disturbance. The central bank observes a private signal and a public nowcast:

$$s_c = \theta + u_c, \qquad u_c \sim N(0, \tau_c^{-1}), \tag{4}$$

$$b = \theta + u_b, \qquad u_b \sim N(0, \tau_b^{-1}). \tag{5}$$

The variable b is the big-data nowcast. Its precision, τ_b, is the theoretical counterpart of satellite nowcast validity. In the empirical implementation, τ_b is higher when cloud cover is lower.

Private agents observe b and idiosyncratic private signals

$$s_j = \theta + u_j, \qquad u_j \sim N(0, \tau_p^{-1}), \quad j \in [0, 1]. \tag{6}$$

This information structure captures two realistic features of monetary-policy environments.

First, the central bank may have an institutional information advantage because it observes supervisory information, internal staff forecasts, and policy deliberations that are not fully public. Second, private agents are not uninformed. They observe public data, produce proprietary forecasts, and purchase information from data providers. The key friction is that these sources differ in precision and timing. A satellite nowcast increases the precision of the public component of the information set before the policy announcement.

3.2 Central-bank inference

Given Gaussian signals, the central bank's posterior mean is

$$E^{CB}[\theta] = \frac{\tau_c s_c + \tau_b b}{\tau_0 + \tau_c + \tau_b}. \tag{7}$$

Substituting (7) into the policy rule gives

$$i = \psi \frac{\tau_c s_c + \tau_b b}{\tau_0 + \tau_c + \tau_b} + \varepsilon_i. \tag{8}$$

Equation (8) is the source of the information effect. The policy rate is an instrument, but it is also a noisy signal of what the central bank has learned about θ.

This point is central for the interpretation of high-frequency monetary-policy surprises. If the policy rule is known, a rate decision can be inverted to infer something about the central bank's assessment. The degree of inference depends on the relative precision of the central bank's private signal and the precision of information already available to the market. When market information is poor, a given rate move carries a large informational surprise. When market information is precise, the same rate move is less likely to be interpreted as hidden macroeconomic news.

3.3 Private inference before and after policy

Before the policy announcement, private agents form expectations using their own signals and the public nowcast. To allow for the fact that market participants care about common interpretations of macroeconomic news, suppose agent j chooses an expectation a_j to minimize

$$E\left[(a_j - \theta)^2 + r(a_j - \bar{a})^2 \mid s_j, b\right], \tag{9}$$

where $r \geq 0$ captures a coordination motive. The pre-announcement expectation is

$$E_j^-[\theta] = \omega_p(r) s_j + \omega_b(r) b, \tag{10}$$

with $\omega_b'(r) > 0$ and $\omega_p'(r) < 0$. A stronger coordination motive raises the effective weight on public information.

The coordination term is not needed for the basic information-effect result, but it is useful in a central-bank communication environment. Financial-market participants often care not only about their own estimate of fundamentals but also about how other participants will interpret common signals. Public forecasts, including widely distributed nowcasts, can therefore receive more attention than their private statistical precision alone would imply. This feature strengthens the role of common information in expectations formation and makes the quality of public nowcasts relevant for asset-pricing responses.

After observing the policy rate, agents can remove the component of i that is mechanically attributable to the already observed nowcast b. Define

$$q \equiv \frac{\tau_0 + \tau_c + \tau_b}{\psi \tau_c} \left(i - \psi \frac{\tau_b}{\tau_0 + \tau_c + \tau_b} b \right) = s_c + \eta_i, \tag{11}$$

where $\eta_i = \frac{\tau_0+\tau_c+\tau_b}{\psi\tau_c}\varepsilon_i$. The precision of q is

$$\tau_q = \left[\tau_c^{-1} + \left(\frac{\tau_0 + \tau_c + \tau_b}{\psi \tau_c} \right)^2 \tau_i^{-1} \right]^{-1}. \tag{12}$$

The post-announcement posterior is

$$E_j^+[\theta] = \frac{\tau_p s_j + \tau_b b + \tau_q q}{\tau_0 + \tau_p + \tau_b + \tau_q}. \tag{13}$$

The posterior in (13) makes clear why a more precise public nowcast can reduce the incremental information content of the policy rate. The weight on q is decreasing in the precision of other signals, holding τ_q fixed. Moreover, the precision of q itself depends on τ_b because the policy rule loads on both the central-bank private signal and the public nowcast. If the public nowcast is already highly informative, extracting additional information from the residual component of the policy rate becomes less important for private beliefs.

3.4 Information effects and the role of nowcast precision

Let $\Delta i_j = i - E_j^-[i]$ denote the policy surprise. The effect of the surprise on expected output decomposes into an information component and a direct policy component:

$$\frac{\partial E^+[y]}{\partial \Delta i} = \underbrace{\frac{\partial E^+[\theta]}{\partial \Delta i}}_{\text{central-bank information effect}} - \underbrace{\phi}_{\text{direct interest-rate effect}}. \tag{14}$$

The information component is proportional to the posterior weight placed on the rate-implied signal q:

$$\frac{\partial E^{+}[\theta]}{\partial \Delta i} \propto \frac{\tau_q}{\tau_0 + \tau_p + \tau_b + \tau_q} \cdot \frac{\partial q}{\partial i}. \tag{15}$$

Proposition 1 (Nowcast precision and policy transmission)**.** *An increase in nowcast precision τ_b reduces the private sector's reliance on the policy rate as a signal of fundamentals once the public nowcast is sufficiently informative. In that region, higher τ_b attenuates the central-bank information effect and increases the relative importance of the direct interest-rate channel.*

The intuition is straightforward. If agents have little independent information about the current state of the economy, policy surprises are interpreted partly as revelations of the central bank's private assessment. If agents have a precise public nowcast, policy surprises are less informative about fundamentals and more cleanly interpreted as policy shocks. The model therefore predicts that valid nowcasts should reduce equity-market information effects and increase bond-market pass-through.

The proposition also clarifies why the sign of an asset-price response is not by itself a sufficient statistic for monetary-policy effectiveness. A positive equity response to a tightening shock can arise when the information effect dominates the conventional discount-rate effect. Such a response does not necessarily mean that monetary policy is expansionary; it means that the announcement has changed beliefs about fundamentals. The empirical implication is that better nowcasting should make equity responses more consistent with the conventional liquidity channel while making bond-yield responses more responsive to the policy component of the shock.

3.5 Testable implications

The model delivers three implications that map directly into the empirical design. First, the precision of the nowcast should matter more for asset prices around policy announcements than on ordinary days. This implication motivates the use of high-frequency monetary-policy surprises and their interaction with BDVI. If BDVI were merely a proxy for the business cycle, it would predict asset prices more generally; the model instead predicts that BDVI should change the interpretation of policy news.

Second, the sign of the BDVI interaction should differ across financial outcomes. For equities, a tightening shock has two effects. The direct discount-rate channel lowers equity prices, while a favorable central-bank information effect can raise them. A more precise nowcast reduces the second component, so the equity response to a tightening should become more negative when BDVI is high. For bond yields, the prediction is different. If policy surprises are interpreted more cleanly as policy shocks, short- and long-maturity yields should respond more strongly to the policy component of the announcement. The model therefore predicts a positive BDVI interaction for yield responses.

Third, the interaction should be stronger when the value of alternative information is higher. Periods of heightened uncertainty, natural disasters, or disruptions to conventional information production should raise the marginal value of a remote-sensing nowcast. This implication motivates the state-dependent local projections in Section 6. The state-dependence exercise is important because it distinguishes the information mechanism from a generic correlation between satellite observability and normal economic conditions.

3.6 Connection to monetary-policy effectiveness

The word effectiveness is used in a specific sense. A policy instrument is effective when a given policy innovation produces the intended change in financial conditions and expectations with limited contamination from unintended inference about fundamentals. In the model, the policy rate combines two objects: a stance component and a signal component. If the private sector cannot separate these components, the same observed rate change can produce ambiguous asset-price responses. A tightening may loosen perceived financial conditions if markets interpret it as good news about the real economy. Conversely, an easing may fail to support expectations if markets infer that the central bank has adverse private information.

Nowcasting improves policy effectiveness by reducing this ambiguity. It does not mechanically make the policy multiplier larger in every state. Rather, it improves the mapping from the policy action to private beliefs. In Bayesian terms, a more precise public nowcast increases the posterior weight placed on public information and reduces the marginal information content of the policy-rate residual. In financial-market terms, the announcement becomes less of a macroeconomic signal and more of a policy signal. In institutional terms, the central bank operates in a better shared information environment.

This distinction is useful for interpreting the empirical results. Stronger pass-through to bond yields under high BDVI is evidence that the policy component is transmitted more cleanly along the yield curve. A more negative equity response to tightening under high BDVI is not evidence that policy is more harmful to equities in a welfare sense; it is evidence that the equity response is less offset by positive information effects. The two outcomes are therefore mutually reinforcing evidence for the same mechanism.

4 Satellite Nowcasting

4.1 Data

The satellite dataset combines six types of remotely sensed information: nighttime lights from NPP/VIIRS, day and night land surface temperature from MODIS, and CO, NO_2, and SO_2 from Sentinel-5P/TROPOMI. These variables capture complementary physical margins of production: light, heat, and emissions. The sample is organized around listed manufacturing firms with identifiable production locations and sufficient history for rolling prediction.

The construction of the firm-level panel follows a micro-to-macro logic. I first identify manufacturing firms for which production locations can be matched to satellite pixels. For each location, satellite observations are extracted at the relevant overpass time and then aggregated to the firm-month level. The aggregation is performed before the macroeconomic prediction exercise so that the object being evaluated is not a single image, but a repeated, firm-level economic signal. This distinction is important. A single satellite observation may be noisy because of atmospheric conditions, local weather, or sensor limitations. A panel of repeated observations across firms and months can still contain stable information about production.

The feature set is designed to capture several dimensions of activity. Level features measure the current intensity of lights, heat, or emissions. Growth features compare current values with the firm's own historical baseline. Volatility features capture instability in the physical signal, which may reflect intermittent production. Extreme-value indicators identify whether the current observation is close to a local high or low relative to the firm's recent history. This feature construction intentionally uses within-firm variation, because the cross-sectional level of satellite intensity may reflect persistent differences in plant size, technology, geography, and local infrastructure rather than changes in current output.

Cloud cover enters the analysis in two ways. In the nowcasting section, it is a diagnostic for whether the satellite signal is operating through genuine visual or thermal observability. In the monetary-policy section, it becomes a source of variation in nowcast validity. The identifying assumption is not that weather is unrelated to the economy in every possible sense. Rather, after controlling for macroeconomic conditions and fixed effects, short-run variation in the clarity of satellite observation affects financial-market responses primarily by changing the precision of the nowcast. This is why the empirical design focuses on interaction effects between monetary-policy shocks and satellite observability rather than treating cloud cover as a direct macroeconomic shock.

4.2 Image extraction and preprocessing

The remote-sensing component is built from repeated image extraction rather than from a single aggregate satellite index. Each listed manufacturing firm is first matched to one or more production locations using firm registration information, subsidiary information, and address-based geographic coordinates. For a firm i with locations $\ell = 1, \ldots, L_i$, I define a local extraction window around each coordinate. The window is large enough to capture the plant and its immediate operating environment, but narrow enough to avoid averaging over unrelated urban or industrial activity. The raw image value for satellite product k is then summarized within the window:

$$z_{i\ell t}^{k} = \frac{1}{|\mathcal{P}_{i\ell}|} \sum_{p \in \mathcal{P}_{i\ell}} z_{pt}^{k} \mathbf{1}\{Q_{pt}^{k} = 1\}, \tag{16}$$

where $\mathcal{P}_{i\ell}$ is the set of pixels assigned to location ℓ, z_{pt}^{k} is the raw pixel value, and Q_{pt}^{k} is a product-specific quality indicator. For optical and thermal products, Q_{pt}^{k} removes observations with poor quality flags or excessive cloud contamination. For atmospheric products, it removes retrievals that fail the quality-assurance screen.

The location-level observations are collapsed to the firm-month level in three steps. First, daily or sub-monthly retrievals are converted into monthly composites using quality-weighted averages. Second, extreme outliers are winsorized within product to reduce the influence of sensor artifacts, fire events, or abnormal retrieval errors. Third, location-level signals are aggregated to the firm level:

$$z_{it}^{k} = \sum_{\ell=1}^{L_i} \omega_{i\ell} z_{i\ell t}^{k}, \qquad \sum_{\ell=1}^{L_i} \omega_{i\ell} = 1. \tag{17}$$

The baseline uses equal location weights when detailed plant-level output weights are unavailable. In robustness checks, firm-level aggregation can be modified by assigning larger weights to headquarters, major subsidiaries, or production locations with stronger historical satellite intensity.

The scale of the data-processing task is substantial. The combined nowcasting model identifies 1,346 satellite-observable firms with out-of-sample accuracy above 60 percent. Over the 2015–2023 sample, a monthly panel contains roughly 108 time observations. This implies approximately 145,000 firm-month observations before product expansion. With six satellite products, the baseline remote-sensing panel contains approximately 872,000 firm-product-month observations before lag construction, rolling moments, and interaction features. Even under a much more conservative count that treats each firm-month composite as a single image chip, the processed imagery is comfortably above 100,000 firm-month image chips. This scale is important because the nowcast is based on repeated, high-dimensional physical measurements rather than on a small hand-collected sample.

Table 1: Approximate scale of the satellite-processing panel

Processing object	Count	Calculation	Interpretation
Satellite-observable firms	1,346	From combined classifier	Firms with prediction accuracy above 60 percent
Monthly sample length	108	2015–2023	Approximate number of firm-month periods
Firm-month image chips	145,368	1,346 × 108	Conservative count before product expansion
Firm-product-month cells	872,208	1,346 × 108 × 6	Baseline panel before lags and rolling features
Highly observable firms	941	From combined classifier	Firms with prediction accuracy above 90 percent

Notes: Counts are approximate and are intended to describe the order of magnitude of the remote-sensing processing task. The realized number of raw image retrievals can be larger because some satellite products have daily or sub-monthly observations before monthly compositing, and because firms may have multiple matched locations.

Table 2: Satellite inputs

Variable	Satellite/product	Resolution	Economic content
Nighttime lights	NPP/VIIRS DNB	about 500 m	Nighttime economic activity
Day land surface temperature	Terra/MODIS MOD11A1	1 km	Industrial heat emissions
Night land surface temperature	Aqua/MODIS MYD11A1	1 km	Operating intensity
CO	Sentinel-5P/TROPOMI	1.2 km	Combustion and industrial emissions
NO_2	Sentinel-5P/TROPOMI	1.2 km	Industrial and transport activity
SO_2	Sentinel-5P/TROPOMI	1.2 km	Heavy industry and energy use

Notes: The table reports the satellite products used to construct the firm-level and aggregate nowcasting signals. Resolution refers to the approximate native spatial resolution of the product used in the source manuscript. Each raw signal is converted into levels, growth rates, rolling moments, and within-firm standardized features before entering the prediction model.

4.3 Feature engineering and validation design

The transformation from satellite imagery to economic predictors is deliberately conservative. The objective is not to fit a complex image-recognition model that is difficult to interpret, but to construct a transparent set of remote-sensing statistics that can be linked to economic activity. For each firm-product-month observation, the pipeline produces a family of features:

$$\text{Level}_{it}^{k} = z_{it}^{k}, \tag{18}$$

$$\text{Growth}_{it,j}^{k} = z_{it}^{k} - z_{i,t-j}^{k}, \qquad j \in \{1, 12, 24, 36\}, \tag{19}$$

$$\text{Momentum}_{it,w}^{k} = z_{it}^{k} - \frac{1}{w}\sum_{r=1}^{w} z_{i,t-r}^{k}, \qquad w \in \{3, 6, 12\}, \tag{20}$$

$$\text{Volatility}_{it,w}^{k} = \left(\frac{1}{w-1}\sum_{r=1}^{w}(z_{i,t-r}^{k} - \bar{z}_{it,w}^{k})^2\right)^{1/2}. \tag{21}$$

The growth features compare current activity with both short-run and year-on-year baselines. The momentum features detect whether the plant is operating above or below its recent normal level. The volatility features capture unstable or intermittent production patterns that are not visible in simple averages.

Because satellite products are measured in different physical units, all predictors are standardized before entering the model:

$$\widetilde{x}_{it}^{k,m} = \frac{x_{it}^{k,m} - \mu_{i,t}^{k,m}}{\sigma_{i,t}^{k,m}}, \tag{22}$$

where m indexes the feature family and the mean and standard deviation are computed using only information available before month t. This real-time standardization avoids look-ahead bias and makes the penalization in the ridge step comparable across products. Quality indicators, missing-value indicators, and cloud-related measures are retained as separate controls rather than silently dropping all imperfect images, since the absence or degradation of an observation can itself be informative about nowcast reliability.

Table 3: Remote-sensing feature families

Feature family	Construction	Economic interpretation
Levels	Current monthly composite of lights, heat, or emissions	Scale of current physical activity
Short-run growth	Month-on-month and recent-window changes	Turning points in production intensity
Year-on-year growth	Difference relative to 12-, 24-, and 36-month baselines	Seasonal-adjusted changes in firm operation
Momentum	Deviation from rolling means	Whether activity is above or below recent normal conditions
Volatility	Rolling standard deviation of product-specific signals	Intermittency, shutdowns, or unstable operation
Extremes and quality flags	Local maxima/minima, missingness, and cloud-quality indicators	Reliability of the image-derived signal

Notes: The table describes the main feature families used to convert raw satellite observations into predictors. All transformations are computed recursively using information available before the prediction month.

The validation design mirrors the real-time problem faced by policymakers. Prediction models are estimated on rolling historical windows and evaluated out of sample. At no point is future revenue or future macroeconomic information used to construct the features for month t. This timing discipline is central for the policy interpretation of the exercise. A nowcast that performs well only after using revised or future information would not improve the information set available at the time of a policy decision.

4.4 Firm-level prediction

For firm i in month t, let X_{it} be the vector of satellite features extracted around the firm's production locations. The target is whether revenue increases:

$$Y_{it} = \mathbf{1}\{\Delta \log(\text{Revenue}_{it}) > 0\}. \tag{23}$$

The prediction model uses rolling ridge regularization followed by a logistic link:

$$\widehat{\beta}_{it} = \arg\min_{\beta} \sum_{\tau=t-35}^{t} (Y_{i\tau} - X_{i\tau}'\beta)^2 + \lambda\|\beta\|_2^2, \tag{24}$$

$$\Pr(Y_{it} = 1 \mid X_{it}) = \frac{\exp(X_{it}'\widehat{\beta}_{it})}{1 + \exp(X_{it}'\widehat{\beta}_{it})}. \tag{25}$$

The rolling window allows the mapping from physical signals to activity to adapt to changes in firm behavior and macroeconomic conditions.

Table 4: Firm-level nowcasting performance

	NTL	LSTD	LSTN	NO_2	CO	SO_2	Combined
Accuracy	59.1%	59.0%	59.3%	57.7%	58.8%	57.5%	81.9%
F1 score	69.0%	69.6%	69.7%	67.7%	70.1%	68.6%	90.3%
Observable firms, accuracy $> 60\%$	749	736	638	940	960	933	1,346
Highly observable firms, accuracy $> 90\%$	545	367	270	551	250	146	941

Notes: The table reports out-of-sample classification performance for the direction of firm revenue growth. The combined column pools nighttime lights, land-surface temperature, and pollution-related satellite features. "Observable firms" count firms for which the rolling prediction accuracy exceeds the threshold shown in the row label.

The combined signal performs substantially better than any single input. This is consistent with the interpretation that different satellite products capture different aspects of production and that aggregation diversifies idiosyncratic measurement error.

Several features of the results are worth emphasizing. First, the individual satellite inputs have modest but nontrivial predictive power. This is expected because each variable captures only one physical dimension of activity. Nighttime lights may be informative for plants with visible illumination, while thermal signals may better capture heat-intensive production, and pollution measures may be more informative for heavy industry. Second, the combined signal performs much better than the individual inputs. This result is consistent with the idea that the economic information is distributed across complementary physical channels. Third, the number of highly observable firms is large enough to support aggregate inference. The combined model identifies 1,346 firms with accuracy above 60 percent and 941 firms with accuracy above 90 percent, suggesting that the nowcast is not driven by a small set of special cases.

The firm-level exercise also disciplines the interpretation of the aggregate results. If the satellite series only correlated with aggregate activity after aggregation, it would be difficult to know whether the signal captured production or merely comoved with broad time trends. The firm-level prediction task is more demanding because it asks whether location-specific physical signals predict firm-specific revenue growth. Passing this test gives the aggregate nowcast a microeconomic foundation.

4.5 Aggregate nowcasting

At the macro level, the analysis tests whether satellite-observable firms represent aggregate activity and whether satellite signals predict PMI. Let Y_t denote a low-frequency macroeconomic outcome and $x_{t-j/m}$ a higher-frequency satellite signal. The MIDAS specification is

$$Y_t = \alpha + \rho Y_{t-1} + \theta \sum_{j=0}^{K} w(j; \gamma_1, \gamma_2) x_{t-j/m} + u_t, \tag{26}$$

where

$$w(j; \gamma_1, \gamma_2) = \frac{\exp(\gamma_1 j + \gamma_2 j^2)}{\sum_{\ell=0}^{K} \exp(\gamma_1 \ell + \gamma_2 \ell^2)}. \tag{27}$$

The MIDAS framework is appropriate because it preserves the high-frequency timing of the satellite signal. A simpler monthly average would discard information about when within the month the signal is observed. The estimated lag weights allow the data to determine whether recent observations or earlier observations have greater predictive content. In a production setting, this flexibility matters because satellite measures may lead official statistics by different amounts depending on the variable. Heat and emissions can move contemporaneously with production, while accounting-based revenue and official activity indicators may be recorded or published later.

Table 5: Aggregate representativeness of satellite-observable firms

	NTL	LSTD	LSTN	NO_2	CO	SO_2	Combined
Correlation with PMI	0.90	0.92	0.92	0.92	0.91	0.91	0.91
p-value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Correlation with GDP	0.94	0.95	0.95	0.94	0.94	0.94	0.94
p-value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Notes: The table reports simple correlations between aggregate indicators and signals constructed from satellite-observable firms. The exercise evaluates whether the satellite-observable firm panel is representative of broader macroeconomic activity before it is used in the policy-transmission analysis.

Table 6: PMI expansion nowcasting

	NTL	LSTD	LSTN	NO_2	CO	SO_2	Combined
Overall accuracy	58.8%	65.0%	58.8%	66.7%	51.7%	75.9%	71.3%
Overall F1	72.0%	78.8%	74.1%	75.0%	61.1%	76.2%	80.6%
Low-cloud accuracy	62.1%	76.1%	71.4%	69.0%	53.3%	81.7%	80.0%
Low-cloud F1	76.6%	86.4%	83.3%	81.6%	69.6%	86.3%	89.9%
High-cloud accuracy	53.3%	54.5%	53.3%	60.0%	50.0%	60.0%	60.3%
Rank-sum p-value	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)

Notes: The table reports classification performance for whether manufacturing PMI is above the expansion threshold. Low-cloud observations are those in which satellite observability is relatively high. The rank-sum test compares the distribution of satellite signals across PMI states.

Table 7: MIDAS estimates using satellite thermal indicators

	Equal-weight LST		Revenue-weighted LST	
	Industrial value added	PMI	Industrial value added	PMI
L^2 LST	0.028*	0.029*	0.023**	0.023**
	(0.011)	(0.011)	(0.010)	(0.010)
L^1 LST	0.012	-0.006	0.008	-0.005
	(0.013)	(0.009)	(0.013)	(0.009)
L^0 LST	-0.012	-0.002		
	(0.008)	(0.006)		

Notes: Standard errors are in parentheses. LST denotes land-surface temperature. L^j denotes the MIDAS lag term. Significance follows the source manuscript. The table is retained as a compact representation of the source estimates and should be replaced by final estimation output in the submission version.

5 Empirical Strategy for Monetary-Policy Transmission

5.1 Nowcast validity

The empirical counterpart of τ_b is a Big Data Validity Index (BDVI) derived from satellite observability. Cloud-cover values are extracted for the headquarters and subsidiaries of satellite-observable manufacturing firms. These values are aggregated across locations, satellite products, and firms:

$$CM_{i,s,t} = \frac{1}{L_{is}} \sum_{\ell=1}^{L_{is}} CM_{i,s,\ell,t}, \tag{28}$$

$$CM_{i,t} = \frac{1}{S_i} \sum_{s=1}^{S_i} CM_{i,s,t}, \tag{29}$$

$$BDVI_t = 1 - \frac{1}{N} \sum_{i=1}^{N} \widetilde{CM}_{i,t}. \tag{30}$$

A higher BDVI denotes clearer satellite observability and a more informative nowcast.

The use of BDVI has two advantages. First, it is measured before the financial-market response to monetary policy, so it is not mechanically generated by asset prices. Second, it reflects a technological constraint on information production rather than an endogenous choice by investors or the central bank. Market participants may choose whether to use satellite information, but they do not choose whether a satellite image is obscured by clouds at the relevant observation point. This feature makes BDVI a useful proxy for the precision of the public nowcast in the model.

The empirical interpretation of BDVI is intentionally narrow. It does not measure the level of economic activity; it measures the expected informativeness of a satellite-based signal about activity. A high BDVI month is a month in which satellite images are clearer and, all else equal, more useful for nowcasting. A low BDVI month is a month in which satellite information is degraded. This distinction helps separate the information channel from the real activity channel.

5.2 Monetary-policy shocks and local projections

The monetary-policy analysis uses daily Chinese financial-market data from 2015 to 2023. Monetary-policy shocks are identified using one-year interest-rate swap quotes linked to the seven-day repo fixing rate around policy events. The main policy-rate proxy is the daily change in DR007. Financial outcomes are CSI 300 returns, three-month Treasury yields, and ten-year Treasury yields.

The use of interest-rate swaps follows the logic of high-frequency identification. Around narrow policy windows, changes in derivative prices are more likely to reflect revisions in expected policy rather than slow-moving macroeconomic conditions. The one-year swap rate is particularly useful because it summarizes expectations about the near-term path of short rates while remaining closely linked to money-market conditions. The empirical specification treats swap surprises as instruments for movements in the policy-rate proxy and for the interaction between policy shocks and BDVI.

Table 8: Descriptive statistics

Variable	Definition	Obs.	Mean	Std. dev.	Min	Max
DR007	DR007 rate (%)	2,328	2.2137	0.4268	1.1379	3.1818
ΔDR007	Daily change in DR007	2,328	-0.0002	0.1145	-0.9246	0.8831
Stock return	CSI 300 daily return (%)	2,328	0.0219	1.1487	-7.0043	7.1689
3m Treasury	3-month Treasury yield (%)	2,328	2.1175	0.5704	0.7819	4.0210
Δ3m Treasury	Daily change, 3m yield	2,328	-0.0003	0.0407	-0.2550	0.3485
10y Treasury	10-year Treasury yield (%)	2,328	2.9261	0.5212	1.5958	3.9851
Δ10y Treasury	Daily change, 10y yield	2,328	-0.0005	0.0225	-0.1701	0.1129
Cloud mask	Satellite cloud-cover index	1,814	0.9235	0.0951	0.7122	1.1783
CPI yoy	CPI inflation (%)	2,328	1.5057	1.2326	-0.8000	5.3800
FCI cycle	Financial conditions cycle	2,328	0.4117	0.1536	0.0238	0.7903
GDP yoy	Real GDP growth (%)	2,328	5.5977	3.2369	-6.8000	18.9000

Notes: The sample covers daily observations from 2015 to 2023. DR007 is the seven-day pledged repo fixing rate. Treasury yields are expressed in percent. The cloud-mask series is the satellite cloud-cover index used to construct BDVI. Macro controls are aligned to daily frequency by carrying forward the most recently available official release.

For each horizon $h = 0, \ldots, 20$, define the cumulative response of outcome Y as

$$R^Y_{t,h} = Y_{t+h} - Y_{t-1}. \tag{31}$$

The baseline local-projection specification is

$$R^Y_{t,h} = \alpha_h + \beta_h \Delta DR007_t + \gamma_h (\Delta DR007_t \times BDVI_t) + \delta_h BDVI_t + X'_{t-1}\Pi_h + \mu_y + \mu_m + \mu_w + \varepsilon_{t,h}. \tag{32}$$

The model is estimated by IV. The first stage instruments $\Delta DR007_t$ with the high-frequency swap surprise ΔIRS_t and instruments the interaction with $\Delta IRS_t \times BDVI_t$. The conditional impulse response is

$$\frac{\partial R^Y_{t,h}}{\partial \Delta DR007_t} = \beta_h + \gamma_h BDVI_t. \tag{33}$$

The coefficient γ_h is the key parameter. The model predicts $\gamma_h < 0$ for equity returns if nowcasts attenuate information effects, and $\gamma_h > 0$ for bond yields if nowcasts strengthen interest-rate pass-through.

The specification includes lagged macro-financial controls and calendar fixed effects. The controls absorb variation in inflation, financial conditions, and output growth that could jointly affect policy and asset prices. Year, month, and day-of-week fixed effects absorb slow-moving regime changes and calendar regularities. The local-projection horizon is set to 20 trading days, approximately one calendar month, which is long enough to trace financial-market transmission while short enough to remain close to the policy event.

The identifying variation comes from the interaction of monetary-policy surprises with the information environment. A concern in monetary-policy studies is that policy surprises may be endogenous to the central bank's information set. The IV strategy addresses this by using high-frequency swap movements as instruments for the policy-rate innovation. A separate concern is that the information environment may be endogenous to macroeconomic conditions. BDVI mitigates this concern because it is based on satellite observability. The interaction design further asks whether the same policy shock has different effects when the nowcast is more or less valid.

5.3 Instrument relevance and exclusion

The first-stage relationship is motivated by the institutional link between DR007 and the interest-rate swap market. Let Z_t denote the high-frequency swap surprise and let $Z_t \times BDVI_t$

denote the corresponding interaction instrument. The first-stage system can be written as

$$\Delta DR007_t = a_0 + a_1 Z_t + a_2(Z_t \times BDVI_t) + a_3 BDVI_t + X'_{t-1}A + \eta_t, \tag{34}$$

$$\Delta DR007_t \times BDVI_t = b_0 + b_1 Z_t + b_2(Z_t \times BDVI_t) + b_3 BDVI_t + X'_{t-1}B + \nu_t. \tag{35}$$

The relevance condition requires that swap surprises predict both the policy-rate innovation and its interaction with BDVI. This condition is natural because swap rates are forward-looking prices of expected short-rate paths. The empirical implementation should report first-stage statistics in the appendix, including the joint strength of the excluded instruments.

The exclusion restriction is that, conditional on controls and fixed effects, the high-frequency swap surprise affects the cumulative financial response only through the monetary-policy innovation and its interaction with the information environment. This assumption is standard in high-frequency monetary-policy identification, but the present design adds an information-state dimension. The identifying object is not the unconditional effect of cloud cover or the unconditional effect of policy. It is the difference in the response to the same type of policy surprise across states in which satellite nowcasts are more or less informative.

Two features make the design more credible. First, the BDVI state is predetermined relative to the financial-market response. Satellite observability is measured from physical image quality rather than from asset prices. Second, the specification includes macro-financial controls and calendar fixed effects, so the interaction is not simply comparing recessions with expansions or crisis months with tranquil months. The remaining identifying assumption is that, after these controls, cloud-induced variation in nowcast precision is not itself a direct driver of the narrow-window response to monetary policy.

5.4 Inference

Local projections generate overlapping dependent variables across horizons, so standard errors must account for serial correlation induced by the construction of $R^Y_{t,h}$. The preferred inference uses heteroskedasticity- and autocorrelation-robust standard errors at each horizon. In a full replication package, the appendix should also report robustness to alternative lag lengths and to block bootstrap confidence intervals. Confidence bands in the figures are interpreted horizon by horizon rather than as simultaneous bands over the entire impulse-response path.

The coefficient path $\{\gamma_h\}_{h=0}^{20}$ is the central empirical object. Because the economic interpretation depends on the pattern across horizons, the paper emphasizes impulse-response figures in the main text and reserves horizon-specific coefficient tables for the appendix. This division mirrors the style of many central-bank working papers: the main text reports visual response functions and key economic magnitudes, while detailed coefficient matrices are placed in supplementary tables.

6 Monetary-Policy Transmission Results

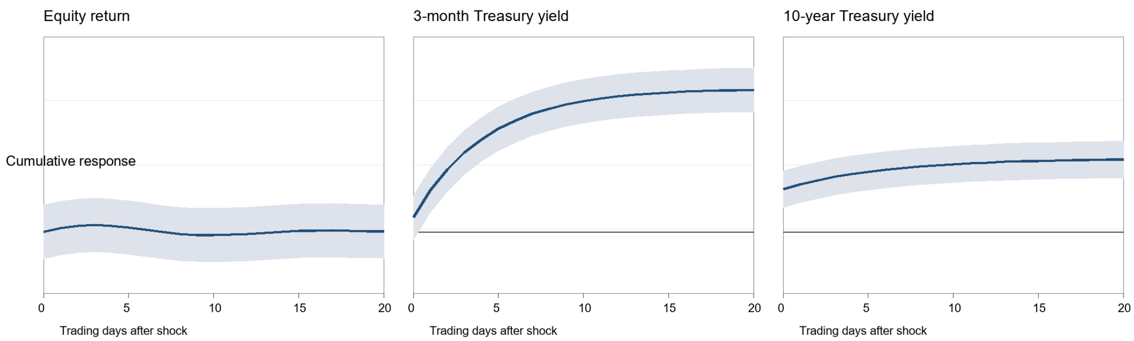

Figure 1: Baseline local-projection responses

Notes: The figure presents an English graphical summary of the baseline impulse-response patterns reported in the source manuscript. Shaded areas denote confidence bands.

The baseline responses show that monetary-policy shocks transmit clearly to bond yields but less cleanly to equities. Short-term yields rise persistently after a tightening shock, and long-term yields also respond positively. Equity returns, however, fluctuate around zero. This pattern is consistent with the presence of an information effect: a tightening can be interpreted not only as a contractionary policy move but also as a signal that the central bank has favorable information about current fundamentals.

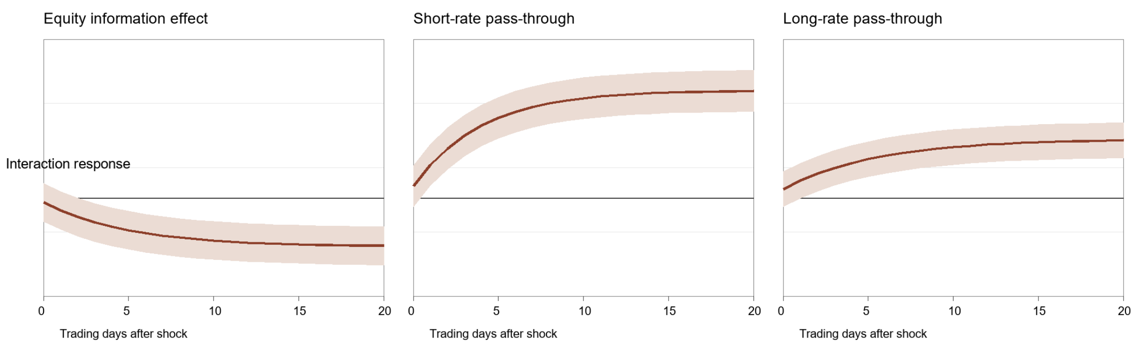

Figure 2: Nowcast validity and monetary-policy transmission

Notes: The figure summarizes the estimated interaction between monetary-policy shocks and satellite nowcast validity. Negative equity responses indicate weaker information-effect distortions; positive bond-yield responses indicate stronger interest-rate pass-through.

The interaction estimates support the mechanism. When satellite observability is higher, the equity-market information effect is smaller: the response of equity returns to tightening shocks becomes more negative, as predicted by the standard liquidity channel. At the same time, the response of short- and long-term bond yields becomes stronger, indicating more effective pass-through along the yield curve. These two results are not contradictory. They are two manifestations of the same information channel. Better nowcasts reduce the extent to which policy surprises are interpreted as hidden macroeconomic news and therefore make the policy component of the announcement more salient.

The magnitude of the interaction is economically meaningful. In low-validity periods, policy surprises contain a larger informational component, and equity prices do not respond in the direction predicted by a simple discount-rate channel. In high-validity periods, the same class of policy surprises generates a cleaner tightening response. For bond yields, higher BDVI raises the sensitivity of both short and long maturities to the policy shock, suggesting that the policy path is incorporated more effectively into the yield curve when the macroeconomic information environment is clearer.

The results also speak to the interpretation of data-dependent monetary policy. A data-dependent policy framework can be misunderstood if the public does not observe the same data or cannot infer how the central bank reads the data. In that case, policy actions may unintentionally become macroeconomic news. Satellite nowcasting reduces this problem by providing a timely external signal of activity. The channel is not that satellite data mechanically cause stronger policy effects. Rather, satellite data improve the common information set, and the improved common information set changes how policy surprises are decoded.

Table 9: Model predictions and empirical patterns

Outcome	Conventional policy channel	Information-effect distortion	High-BDVI prediction
Equity returns	Tightening lowers prices	Favorable information can raise prices	More negative response
Short-term yields	Tightening raises yields	Ambiguous if policy news is mixed	Stronger pass-through
Long-term yields	Tightening raises expected path	Damped by macro-signal confusion	Stronger pass-through
Disaster-period responses	Transmission may be noisy	Information scarcity is larger	Larger BDVI interaction

Notes: The table summarizes the sign restrictions implied by the Bayesian information model and compares them with the empirical patterns reported in the impulse-response analysis. It is a theory-to-evidence map rather than a coefficient table.

6.1 Economic magnitude

The mechanism is economically relevant because it concerns the interpretation of the marginal policy surprise. A one-time improvement in satellite observability does not change the structural interest-rate channel by itself. Instead, it changes the decomposition of an announcement into policy news and macroeconomic news. In low-BDVI states, a tightening surprise may be read as an indication that the central bank has observed stronger activity than the market expected. This interpretation offsets the discount-rate channel in equities and can dampen the response of longer-maturity yields. In high-BDVI states, the public already observes a more precise signal of activity, so the same tightening surprise is less likely to be interpreted as hidden growth news. The announcement therefore has a cleaner contractionary interpretation.

This interpretation also explains why the effect is visible in financial variables before it is visible in real activity. Asset prices aggregate beliefs immediately after a policy event. If nowcast precision changes belief updating, the first place to observe the mechanism is in equity and bond markets. Real activity may respond later, but a real-side response would be mediated through the financial conditions whose high-frequency response is studied here. The paper therefore treats financial-market pass-through as the first stage of monetary-policy effectiveness.

6.2 Relation to central-bank communication

The results are closely related to the literature on central-bank communication. Communication can reduce information effects by explaining the reaction function, the assessment of fundamentals, or the conditional path of policy. Satellite nowcasting operates through a complementary channel. It improves the outside information set against which communication is interpreted. If private agents have a better independent assessment of current conditions, a policy statement needs to carry less of the burden of revealing the state of the economy.

This distinction is useful in a data-dependent regime. The phrase data dependence can be informative only if the public understands the data to which policy is responding. When official data are delayed, the reaction function is harder to infer because the relevant state variable is partly unobserved. A timely satellite nowcast narrows that gap. The implication is not that alternative data replace communication, but that alternative data can make communication more effective by anchoring the common factual basis of the policy discussion.

6.3 State dependence

The information value of nowcasts should be higher when the macroeconomic environment is more uncertain. To test this implication, I estimate a state-dependent local-projection model:

$$R^{Y}_{t,h} = \alpha_h + \sum_{s\in\{0,1\}} \mathbf{1}\{S_{t-1} = s\} \left[\beta_{h,s}\Delta DR007_t + \gamma_{h,s}(\Delta DR007_t \times BDVI_t)\right] + X'_{t-1}\Pi_h + \varepsilon_{t,h}, \tag{36}$$

where S_{t-1} identifies disaster periods using direct economic losses relative to GDP. The source estimates indicate that the BDVI interaction is more than three times larger in disaster periods than in normal periods and remains persistent for longer horizons. The finding is consistent with the model: when conventional information production is impaired, the marginal value of a timely satellite nowcast rises.

The state-dependent results help distinguish the information mechanism from a generic data-correlation story. If satellite observability were merely proxying for normal macroeconomic conditions, the interaction would not necessarily be stronger in states in which official and private information production are disrupted. The fact that the effect is larger in disaster periods is consistent with a marginal-value interpretation: when uncertainty is high and conventional data collection is more difficult, an independent remote-sensing signal becomes more valuable for expectation formation.

7 Robustness, Interpretation, and Limitations

The empirical design is subject to several natural concerns. The first is that cloud cover may be correlated with real economic disruptions rather than only information quality. This concern is partly addressed by the structure of the test. The main coefficient of interest is not the unconditional relationship between cloud cover and asset prices, but the interaction between policy surprises and nowcast validity. The specification also controls for macroeconomic conditions and fixed effects. In addition, the interpretation relies on the fact that cloud cover degrades the quality of satellite observation in a mechanically understood way. The same physical constraint that weakens image quality also weakens the nowcast.

A second concern is that satellite nowcasts may be used by only a subset of sophisticated investors. The mechanism does not require universal direct use of satellite data. It requires that satellite-based information contributes to the public or semi-public information environment. Information can diffuse through data vendors, broker research, analyst forecasts, news coverage, and institutional trading. If some informed agents incorporate satellite signals into forecasts and prices, the common information set faced by other participants can improve even without universal direct access to raw images.

A third concern is that the local-projection estimates summarize financial-market responses rather than real activity responses. This is intentional. The mechanism in the paper concerns the transmission of policy through expectations and asset prices. Financial-market responses are the natural high-frequency objects for identifying information effects, because they react quickly to announcements and are less contaminated by later macroeconomic shocks. The paper does not claim that satellite nowcasting alone determines real outcomes; it argues that nowcasting improves the first stage of the transmission process by making policy signals cleaner.

A fourth concern is that the figures in the current draft summarize impulse-response patterns rather than presenting the full coefficient matrix. This reflects the status of the draft and the available source material. In a submission version, the impulse-response figures should be generated directly from the final estimation files, and a supplementary table should report first-stage statistics, horizon-specific coefficients, and standard errors. The current draft preserves the economic content and structure of the empirical results while leaving room for replacement with final estimation output.

The main limitation is that satellite imagery is not a universal substitute for official statistics. It is most informative for sectors with observable physical footprints, especially manufacturing and infrastructure-related activity. Service-sector activity, digital consumption, and household expectations are less directly visible from satellite images. The policy implication is therefore not to replace official data, but to build a layered information system in which satellite nowcasting complements surveys, administrative data, financial data, and central-bank staff analysis.

7.1 Robustness architecture

A central-bank working-paper version should make the robustness architecture explicit even when some tables are moved to the appendix. The first set of checks should vary the definition of BDVI. The baseline index is continuous, but the mechanism should also appear when BDVI is discretized into high- and low-observability states. Additional versions can use alternative cloud-cover thresholds, product-specific cloud masks, and revenue-weighted aggregation across firms. These checks ask whether the result depends on one particular normalization of satellite observability.

The second set of checks should vary the policy shock. The baseline uses high-frequency swap surprises as instruments for DR007 innovations. A natural robustness exercise is to replace the one-year swap surprise with shorter-maturity swap movements, policy-window changes in repo rates, or an event-day monetary-policy surprise constructed from money-market instruments. The model predicts that the BDVI interaction should be strongest for shocks that markets plausibly interpret as policy news and weaker for movements that are contaminated by liquidity or non-policy factors.

The third set of checks should vary the outcome window. The baseline horizon of 20 trading days is useful because it captures short-run financial transmission without drifting too far from the policy event. Robustness checks at 5, 10, and 30 trading days can show whether the results are immediate, persistent, or transitory. A clean information-channel interpretation predicts that the effect should appear quickly after the policy event, while the persistence of yield responses may depend on how the policy path is repriced.

The fourth set of checks should examine sample composition. Excluding major holidays, extreme weather episodes, and large market-stress days can help address the concern that BDVI captures unusual macro-financial states. Splitting the sample before and after major changes in the monetary-policy framework can also test whether the information channel is stronger when markets place greater weight on price-based policy signals.

7.2 Interpretation for policy design

The results have a narrow but important policy interpretation. A central bank that invests in better nowcasting infrastructure improves more than its internal forecast. It also improves the public environment in which policy is interpreted, provided that nowcast information is made credible, timely, and sufficiently transparent. This does not require the central bank to publish raw satellite data or machine-learning models in real time. It does imply that systematic use of alternative data can support clearer communication about current conditions and reduce the informational ambiguity of policy actions.

The mechanism is particularly relevant for economies in which official data are released with meaningful lags and where financial markets learn from policy actions. In such settings, policy surprises can easily mix stance news and state news. A better real-time measurement system reduces the scope for that mixture. For the conduct of policy, this means that measurement technology is part of the transmission environment. For empirical research, it means that the strength of policy transmission may vary with the precision of the information set available to the public.

8 Conclusion

This paper argues that satellite nowcasting can improve monetary-policy effectiveness by changing the information environment in which policy is interpreted. Satellite imagery provides a timely signal of economic activity before official statistics are released. In a Bayesian model of central-bank information effects, a more precise public nowcast reduces private agents' reliance on policy actions as signals of fundamentals. Empirically, cloud-cover variation in satellite observability shows that better nowcast validity weakens equity-market information effects and strengthens interest-rate pass-through to bond yields.

The broader implication is that nowcasting is not only a forecasting exercise. For a data-dependent central bank, timely and credible measurement technologies can become part of the transmission mechanism itself. By improving the common information set of policymakers and markets, satellite nowcasting can make policy signals cleaner, expectations more aligned, and monetary-policy transmission more effective.

A Technical Appendix

A.1 Derivation of the rate-implied signal

Starting from the policy rule,

$$i = \psi \frac{\tau_c s_c + \tau_b b}{\tau_0 + \tau_c + \tau_b} + \varepsilon_i, \tag{37}$$

subtract the component explained by the public nowcast:

$$i - \psi \frac{\tau_b}{\tau_0 + \tau_c + \tau_b} b = \psi \frac{\tau_c}{\tau_0 + \tau_c + \tau_b} s_c + \varepsilon_i. \tag{38}$$

Multiplying both sides by $(\tau_0 + \tau_c + \tau_b)/(\psi\tau_c)$ yields

$$q = s_c + \frac{\tau_0 + \tau_c + \tau_b}{\psi\tau_c} \varepsilon_i. \tag{39}$$

Thus the rate-implied signal is a noisy observation of the central-bank private signal. Since $s_c = \theta + u_c$, the noise in q is the sum of central-bank signal noise and scaled policy noise. This gives the precision expression in equation (12).

A.2 Interpretation of the nowcast-precision comparative static

The sign of the comparative static with respect to τ_b can be non-monotone because two forces operate simultaneously. A more precise public nowcast makes the central bank's reaction function easier to understand, since the public can better infer which part of the rate decision is a response to commonly observed fundamentals. At the same time, the same public nowcast reduces the private sector's need to infer fundamentals from the residual component of the policy rate. In the empirically relevant region emphasized in this paper, the second force dominates: improvements in satellite nowcast validity reduce the information-effect component of policy surprises.

A.3 Feature construction in the satellite nowcast

For each firm-location pair, the raw satellite variables are converted into a set of standardized features. Let $z^k_{i\ell t}$ denote satellite variable k for firm i, location ℓ, and month t. The feature vector contains the current level, rolling means, growth rates relative to one-, two-, and three-year windows, rolling standard deviations, and indicators for whether the current observation lies near a local maximum or minimum. Location-level values are averaged within firm and then standardized within variable before entering the prediction model. The purpose of standardization is to make signals measured in different physical units comparable in the regularized prediction step.

A.4 Suggested additional tables for the next empirical draft

The current Overleaf draft is structured so that final estimation output can be inserted without changing the paper architecture. For a complete submission version, the following additional tables would be natural:

1. First-stage regressions for the LP-IV estimates, including Kleibergen-Paap or first-stage F statistics.
2. Horizon-specific interaction coefficients for the equity, short-yield, and long-yield responses.
3. Robustness checks using alternative BDVI definitions, including high-cloud indicators and alternative aggregation weights.
4. Robustness checks excluding periods of severe weather, major holidays, and extreme market stress.
5. A table comparing satellite nowcasts with analyst-consensus forecasts and simple autoregressive benchmarks.

These additions would move the paper closer to a full central-bank working-paper submission while keeping the main text focused on the mechanism.

B Additional Derivations

B.1 Posterior updating with two public signals

The baseline model can be generalized to allow both a conventional public signal p and a satellite nowcast b. Let

$$p = \theta + u_p, \qquad u_p \sim N(0, \tau_p^{-1}), \tag{40}$$

$$b = \theta + u_b, \qquad u_b \sim N(0, \tau_b^{-1}), \tag{41}$$

$$s_c = \theta + u_c, \qquad u_c \sim N(0, \tau_c^{-1}). \tag{42}$$

With prior $\theta \sim N(\bar{\theta}_0, \tau_0^{-1})$, the central bank's posterior mean before setting policy is

$$m_c = \frac{\tau_0 \bar{\theta}_0 + \tau_p p + \tau_b b + \tau_c s_c}{\tau_0 + \tau_p + \tau_b + \tau_c}. \tag{43}$$

Suppose the policy rule is

$$i = \psi m_c + \varepsilon_i, \qquad \varepsilon_i \sim N(0, \tau_i^{-1}). \tag{44}$$

Private agents observe p, b, and i, but not s_c or ε_i. Conditional on p and b, the policy rate is a noisy signal of s_c. Rearranging the policy rule yields

$$\frac{\tau_0 + \tau_p + \tau_b + \tau_c}{\psi \tau_c} i - \frac{\tau_0 \bar{\theta}_0 + \tau_p p + \tau_b b}{\tau_c} = s_c + \frac{\tau_0 + \tau_p + \tau_b + \tau_c}{\psi \tau_c} \varepsilon_i. \tag{45}$$

The left-hand side is the rate-implied signal. Its precision falls when policy noise is high and rises when the central bank's private signal receives a larger weight in the policy rule. The posterior of private agents after observing the policy rate is a precision-weighted average of the prior, the two public signals, and the rate-implied signal. This formulation shows that the policy rate is informative only because it is correlated with the central bank's private signal.

B.2 Comparative static

Let ω_q be the posterior weight placed on the rate-implied signal q:

$$\omega_q(\tau_b) = \frac{\tau_q(\tau_b)}{\tau_0 + \tau_p + \tau_b + \tau_q(\tau_b)}. \tag{46}$$

The dependence of τ_q on τ_b reflects the fact that the policy rule itself places weight on the public nowcast. Differentiating gives

$$\frac{\partial \omega_q}{\partial \tau_b} = \frac{\tau_q'(\tau_0 + \tau_p + \tau_b + \tau_q) - \tau_q(1 + \tau_q')}{(\tau_0 + \tau_p + \tau_b + \tau_q)^2}. \tag{47}$$

The expression has two terms. The first captures the possibility that a better nowcast makes the rate-implied signal easier to interpret because the public component of the rule is better measured. The second captures the dilution effect: as τ_b rises, private agents already have a more precise public signal of fundamentals and therefore place less incremental weight on the policy rate. The empirical hypothesis in the paper is that the dilution effect dominates in the relevant region, so the information component of policy surprises declines with nowcast precision.

C Additional Empirical Specifications

C.1 Binary nowcast-validity specification

An alternative to the continuous interaction in equation (32) is a binary high-validity specification:

$$R^Y_{t,h} = \alpha_h + \beta_h \Delta DR007_t + \gamma_h(\Delta DR007_t \times HighBDVI_t) + \delta_h HighBDVI_t + X'_{t-1}\Pi_h + \mu_y + \mu_m + \mu_w + \varepsilon_{t,h}. \tag{48}$$

Here $HighBDVI_t$ equals one when satellite observability is above a pre-specified percentile threshold. This specification is easier to interpret because γ_h compares policy transmission in high- and low-validity states. The continuous specification is preferred in the main text because it uses more variation, but the binary specification is useful as a robustness check.

C.2 Forecast-comparison specification

The nowcasting section can also be expressed as a formal forecast-comparison exercise. Let $\widehat{Y}^{sat}_t$ be the satellite-based forecast and $\widehat{Y}^{bench}_t$ be a benchmark forecast based on lagged official data. The relative loss is

$$d_t = L(Y_t, \widehat{Y}^{bench}_t) - L(Y_t, \widehat{Y}^{sat}_t), \tag{49}$$

where $L(\cdot)$ is either squared error for continuous outcomes or log loss for classification outcomes. A positive mean of d_t indicates that satellite nowcasts improve on the benchmark. In a final empirical version, this comparison should be reported for PMI, industrial value added, and firm-level revenue growth.

C.3 Placebo specification

A placebo test can replace the event-day monetary-policy surprise with non-event-day changes in the same interest-rate instrument:

$$R^Y_{t,h} = \alpha_h + \beta_h \Delta DR007^{placebo}_t + \gamma_h(\Delta DR007^{placebo}_t \times BDVI_t) + X'_{t-1}\Pi_h + \varepsilon_{t,h}. \tag{50}$$

The model predicts that γ_h should be weaker or absent in the placebo sample because the information effect is tied to policy announcements. If BDVI interacted with interest-rate changes for reasons unrelated to monetary-policy information, the placebo interaction would remain large.

References

Paolo Angelini, Gonzalo Camba-Mendez, Domenico Giannone, Lucrezia Reichlin, and Gerhard Rünstler. Interpolation and backdating with a large information set. *Journal of Economic Dynamics and Control*, 35(12):2016–2031, 2011.

Susan Athey and Guido W. Imbens. The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2):3–32, 2017.

Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

Bank for International Settlements. Bis annual economic report 2021. Technical report, Bank for International Settlements, 2021.

Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. Technical report, NBER Working Paper No. 29909, 2022.

Ben S. Bernanke. Monetary policy in a data-rich environment. *Journal of Monetary Economics*, 59(1):1–2, 2012.

Zhicun Bian and Yi Zhang. Monetary policy rules and macroeconomic fluctuations: Evidence from china. *Journal of Financial Research*, (9):28–40, 2012. In Chinese.

Francesco Bianchi and Leonardo Melosi. Dormant shocks and fiscal hegemony. Technical report, Federal Reserve Bank of Chicago Working Paper No. 2013-27B, 2013.

Alan S. Blinder, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4):910–945, 2008.

Brandyn Bok, Daniele Caratelli, Domenico Giannone, Argia Sbordone, and Andrea Tambalotti. Macroeconomic nowcasting and forecasting with big data. Technical report, Federal Reserve Bank of New York Staff Reports, 2018.

Jeffrey R. Campbell. The new keynesian liquidity trap. *Journal of Monetary Economics*, 60 (5):545–560, 2013.

Chuanglian Chen, Tingguo Zheng, and Shujie Yao. Time-varying parameter taylor rule and the stance of central-bank monetary policy. *Economic Research Journal*, (8):43–56, 2016. In Chinese.

Sanjiv Das and Chen Song. Monetary policy surprises and information channels. *Review of Economics and Statistics*, 2023. Forthcoming.

Dave Donaldson and Adam Storeygard. The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–198, 2016.

European Central Bank. Eurosystem report on expert judgment 2021. Technical report, European Central Bank, 2021.

Zhiyong Fan and Yan Li. Expectation management, monetary policy and macroeconomic dynamics in china. *Economic Research Journal*, (5):56–72, 2024. In Chinese.

Domenico Giannone, Lucrezia Reichlin, and David Small. Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4): 665–676, 2008.

J. Vernon Henderson, Adam Storeygard, and David N. Weil. Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028, 2012.

Patrick C. Higgins. *Nowcasting: A Bayesian approach to predict economic indicators using high-frequency data*. PhD thesis, Duke University, 2014.

Marek Jarocinski and Peter Karadi. Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020.

Neal Jean, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353 (6301):790–794, 2016.

Òscar Jordà, Moritz Schularick, and Alan M. Taylor. Betting the house. *Journal of International Economics*, 96(S1):S2–S18, 2015.

Jianhao Lin, Liangyuan Chen, and Yingping Huang. Identifying multiple shocks in monetary policy operations and their economic effects. *Journal of Financial Research*, (5):20–38, 2024. In Chinese.

Viktor Mayer-Schoenberger and Kenneth Cukier. *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt, Boston, 2013.

Alexandre Moreira and Frederique Savignac. Monetary policy transmission in brazil: Evidence from high-frequency identification. *Journal of Monetary Economics*, 2025. Forthcoming.

Sendhil Mullainathan and Jann Spiess. Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106, 2017.

Emi Nakamura and Jon Steinsson. High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*, 133(3):1283–1330, 2018.

Emre Ozturk and Xuguang Simon Sheng. Measuring and modeling ambiguity in health utility analysis. *Journal of Health Economics*, 58:13–27, 2018.

Valerie A. Ramey. Macroeconomic shocks and their propagation. In John B. Taylor and Harald Uhlig, editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier, 2016.

Valerie A. Ramey and Sarah Zubairy. Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, 126(2):850–901, 2018.

Yuhuang Shang and Tingguo Zheng. Mixed-frequency measurement of china's financial conditions index and its warning behavior. *Journal of Financial Research*, (3):21–35, 2018. In Chinese.

James H. Stock and Mark W. Watson. Slack and cyclical unemployment. Technical report, NBER Working Paper No. 24980, 2018.

Richard H. Thaler. Anomalies: The january effect. *Journal of Economic Perspectives*, 1(1): 197–201, 1987.

Xia Wang and Tingguo Zheng. Measuring china's macroeconomic uncertainty based on real-time information flows. *Economic Research Journal*, (10):55–71, 2020. In Chinese.

Xiandong Yan and Wenbo Gao. Central-bank information disclosure and inflation expectation management in china. *Journal of Financial Research*, (8):35–49, 2017. In Chinese.

Chengsi Zhang, Zehao Liu, and Xuezhe Zhu. Micro foundations of monetary-policy transmission channels: Evidence from firm heterogeneity. *Journal of Financial Research*, (3):1–18, 2022. In Chinese.

Tingguo Zheng, Xinyue Fan, and Shiqi Ye. Growth-at-risk measurement and risk sources under macro big data. *Economic Research Journal*, (11):133–152, 2023. In Chinese.

Tingguo Zheng, Xinyue Fan, and Wei Jin. Cpi inflation forecasting and method comparison using macro big data. *Journal of Management Sciences in China*, (8):1–16, 2025. In Chinese.