

Predicting Rare Earth Supply Shocks Using a Machine Learning Classifier Built from Open Geopolitical Data

AYAAN SYED IMAM

Abstract: China is responsible for some 90% of the global processing capacity for rare earth elements (REEs) and can therefore use its export quotas as a geopolitical tool (International Energy Agency, 2024; OECD, 2023). These bans have an impact on markets, as the gaps in supply and demand can be measured, and news of these restrictions is made available in Chinese but not released in English until 48 to 72 hours later on official Chinese outlets (Kuo, 2021). In this paper, I explore the possibility of a machine learning classifier being able to anticipate a REE supply shortage—an event that has yet to be obvious from the market—in all cases for which I can find publicly available data. Three open-source datasets were combined, the GDELT global event database (Leetaru & Schrodt, 2013) with Chinese state media coverage indicators and daily stock prices for two Indian downstream companies and the REMX (rare earth) ETF for 2020–2025 to create 1,556 daily observations. The 10 features were designed and trained in a logistic regression classifier using a strict time-ordered train/test split without data leakage. If applied to unseen data on 2025 the model's recall was 0.87 and its ROC-AUC was 0.732, and the first alert was made an average of 28.5 days before the two real restriction events of 2025, when the REMX ETF was flat or declined. The largest single indicator was the quantity of Chinese state media coverage, not the sentiment. The model triggers about one false alarm every two trading days, showing some limitations in practice. However, the findings show that freely available geopolitical signals are actually predictive in a meaningful way of rare earth price shocks, with lead time.

Keywords: Rare Earth Elements, Geopolitical Risk, Machine Learning, Gdelt, Supply Chain, Export Restrictions, Early Warning Systems, Commodity Markets, China, Critical Minerals

Introduction

Rare earth elements are 17 metals that are embedded in almost all technology these days. Neodymium and dysprosium are needed for permanent magnets in motors for electric vehicles and wind turbines, and gallium and Germanium for semiconductors and defence optics. The IEA considers these materials as key to the clean energy transition and notes a “dangerous concentration” in the supply chains of these materials (International Energy Agency, 2024). Nearly 90% of the world's rare earths are refined in one country, China (OECD, 2023; Zhou et al., 2017). This concentration would give China the power to not only set rare earth prices, but who can buy them. In the past three years, China has proved itself willing to use this leverage. In July 2023 it imposed export controls on gallium and germanium (China State Council, 2023). In April 2025 it deemed seven rare earth elements to be restricted, and in October 2025 extended the restrictions to 12 (China State Council, 2025). All these steps caused a “shock effect” within the companies relying on the materials. The closest financial indicator to this shock is REMX, a U.S.- traded exchange-traded fund that focuses on rare-earth and strategic-metal mining shares. Once a restriction is declared, the price of REMX will normally increase rapidly as investors expect less supply and more expensive prices for the rest of the producers. Indian electronics companies like Bharat Electronics (BEL) and Dixon Technologies are affected by the same shocks because of their dependence on Chinese-made components (Mancheri et al., 2019).

Timing is the key to making this tractable prediction problem. A Chinese export restriction is a policy choice that occurs in the wake of a clear and marked policy signaling process. Notices are published by government ministries, and state-controlled media outlets, such as Xinhua, the Global Times, the Ministry of Commerce (MOFCOM) and People's Daily fill the days and weeks leading up to the publication. This rise is first observed in the Chinese-language information space and takes another 2-3 days to be translated to a sellable headline from the Western financial media (Du et al., 2024; Kuo, 2021). The hypothesis of the present study is that this geopolitical signal is observable in public data and can be learned by a machine learning classifier that alerts when the signal falls within its quiet window, and prior to the announcement and market action.

Three contributions follow from this framing. First, the study assembles a fully open and reproducible daily panel that joins geopolitical event data, a state-media coverage filter and downstream equity prices over a six-year window, so that every input can be reconstructed by another researcher without proprietary terminals or paid news feeds. Second, it treats supply restriction as a forecasting problem with a defined operating threshold and an explicit cost asymmetry, rather than as an after-the-fact event study of announcements that have already occurred. Third, it extends the analysis across a border by including two Indian downstream manufacturers, which allows the transmission of a Chinese policy shock to be traced into an importing economy rather than only into the mining equities that sit closest to the restricted commodity. Taken together, these choices position the paper between the political-economy literature on rare earth statecraft and the quantitative literature on event-driven market prediction, and the following section reviews both. Figure 1 sets out the conceptual pathway

that links supply concentration to the geopolitical leverage and downstream risk examined in the remainder of the paper.




Figure 1. Conceptual pathway linking rare earth supply concentration, geopolitical leverage, and supply-chain risk.

Section 2 summarizes the existing literature on REE geopolitics, event-driven financial prediction and the GDELT database. The data sources, feature engineering and modelling methodology are described in Section 3. The results of sections 4 and 5 are presented and discussed. Conclusions and limitations and directions for further work are given in Section 6.

Literature Review

The literature relevant to this study falls into four strands. The first concerns the political economy of rare earth supply and the use of export policy as an instrument of statecraft. The second examines how rare earth and critical mineral markets price in policy news, largely through event studies and volatility modelling. The third covers the use of large-scale machine-coded news corpora, and GDELT in particular, as a quantitative measure of geopolitical risk. The fourth is methodological and concerns how classifiers of rare, time-ordered events should be built and judged. Each strand is reviewed in turn, and the section closes by identifying the gap that the present study addresses. Scholars have taken notice of the geopolitical factor affecting rare earths supply chain. Mancheri et al. (2019) studied the impact made by China’s policies on the resilience of rare earths supply chain having assessed trade and production statistics. They have concluded that the restrictions imposed on export by China since 2010 have negatively impacted the security of supply for downstream nations. Li et al. (2023) examined the way in which the prices of rare earths change against the influence of political events and economics, thus proving the vital influence of political risk on the price fluctuations for rare earths. Guo et al. (2024) assessed the global rare earth market from the standpoint of competition through grey prediction models. The researchers predict that the growing demand for rare earths will become a serious problem in terms of ensuring diversification of supply.

The historical record behind these findings is well documented. From Chinese rare earth policy across five phases from 1975 to 2018 and show that the objective of the state shifted over time from encouraging upstream extraction, to promoting downstream manufacturing, and finally to consolidating a fragmented industry and curbing unsanctioned production. The importance of that periodisation for the present study is that export instruments have never been used in isolation: quotas, production caps, environmental enforcement and industrial reorganisation have historically been announced in clusters, which is precisely why a coverage-volume signal spanning several ministries may carry more information than any single announcement does.

Barteková and Kemp (2016) approach the same question from the demand side, comparing how the European Union, the United States, Japan and other importing regions responded to supply risk, and conclude that although the stated objectives converged, the chosen instruments diverged sharply according to national policy style and resource endowment. Vekasi (2019) documents the Japanese response to the 2010 restrictions in detail, showing how firms and the state combined stockpiling, substitution research and overseas equity investment into a single strategy. Baskaran and Schwartz (2025) provide the most recent assessment of the 2025 measures and emphasise that the newer controls are licensing-based rather than quota-based, which makes their effect on any given importer harder to observe directly from trade statistics and correspondingly increases the value of indirect indicators of the kind used here. Figure 2 summarises the information pathway that this body of work implies, running from geopolitical events and state-media coverage through to the observable market and supply-chain responses modelled below.




Figure 2. Proposed information pathway from geopolitical events and state-media coverage to observable market and supply-chain responses.

A second strand asks how quickly and how completely markets absorb this kind of policy news. Proelss et al. (2018) examine the rare earth market around the 2012 World Trade Organization dispute and find a structural break in REE prices at the announcement, together with an increase in the stock-price informativeness of REE-sector firms afterwards; they also report that REE prices continued not to follow a random walk, which implies that the market is not fully efficient with respect to policy information. That finding bears directly on the present paper, because a lead time of the kind reported in Section 4 is only possible if some portion of the relevant information remains unpriced for a period. The broader finance literature supplies the mechanism. Tetlock (2007) shows that the content of financial media carries information about subsequent returns rather than merely reflecting them, Bollen et al. (2011) demonstrate that aggregate mood extracted from social text has predictive content for equity indices, and Baker et al. (2016) construct a newspaper-based uncertainty index that anticipates movements in investment and output. What these studies share is the premise that text volume and text tone are measurable economic variables; what they generally lack is application to a commodity whose supply is controlled administratively by a single state.

The investigation into how geopolitics impacts the strategic mineral sector has taken off in recent years. According to Zhang and colleagues (2024), publishing in *Heliyon*, they managed to calculate the degree of risk spillover between geopolitical events and the 5G, semiconductor and rare earth sector of China and discovered that real undertakings such as prohibitions on exports and military actions bring about more significant effects than just perceived risk. Seiler (2025), publishing in *Risks*, determined how the risk seen in the rare earth sector influences other sectors of the economy using quantile regression and found that volatility coming from

the Danish REE market travels downstream to such major sectors as solar energy, transport, and military affairs. Bouoiyour et al. (2026), writing in *Environmental Earth Sciences*, reported on the research they performed on price forecasting of the rare earth market using artificial neural networks with data coming from geopolitics.

Two features of this recent work are worth drawing out. The first is that several of these studies distinguish between realised and perceived risk and consistently find the former to dominate: Zhang et al. (2024) report that concrete actions such as export prohibitions move markets more than generalised risk perception does. If realised policy actions carry the larger effect, then the analytically useful task is to detect the approach of an action before it is realised, which is the task attempted here. The second is that most of this literature is explanatory rather than predictive. Quantile regression, spillover decomposition and grey prediction models are estimated over the full sample, which is appropriate for measuring the size of a relationship but leaves open whether that relationship could have been exploited in advance. Even Bouoiyour et al. (2026), who apply artificial neural networks to REE price forecasting with geopolitical inputs, frame the exercise as forecasting a price series rather than as issuing a discrete, dated warning that a policy event is imminent. The distinction between explaining variance and raising a timely alert is the space that the present study occupies. The database of the Global Database of Events, Language, and Tone or GDELT has been approved for use as a tool for measuring geopolitical risk in literature that has undergone peer evaluation. Leetaru and Schrodt (2013) identified the structure of the GDELT database and also gave an account of the breadth of the database which covers the real-time tracking of the news through television and print news media and even on the Internet in more than 100 languages. Du et al. (2024) used data collected from GDELT to establish the geopolitical risk in the Southeast Asian countries. The research conducted by Ferrara et al. (2020) showed how the indicators got from GDELT can be used in the analysis of financial markets. Caldara and Iacoviello (2022) created the well-known Geopolitical Risk Index (GPR), based on which GDELT-based Goldstein scores have achieved good results (Wang et al., 2025). The given research develops the research by analyzing the information from GDELT to predict the rare-earth elemental supply shock, which has not been studied before.

A fourth strand of literature is methodological and concerns how classifiers of rare, time-ordered events should be constructed and evaluated. Ward et al. (2010) argue, in the context of civil conflict forecasting, that statistical significance is a poor guide to predictive value and that models should be judged on out-of-sample performance instead; their critique applies with full force to a setting in which the positive class is defined by a small number of discrete policy events. Bergmeir and Benítez (2012) analyse the conditions under which cross-validation is and is not admissible for time-series predictors and show that standard random-fold cross-validation breaks the temporal ordering that any realistic deployment must respect, which is the reason the present study uses a single chronological split rather than a shuffled one. Class imbalance introduces a further complication: Chawla et al. (2002) established over-sampling of the minority class as a standard remedy, while the class-weighting approach adopted here achieves a comparable effect without synthesising observations that never occurred. On evaluation, Fawcett (2006) provides the standard treatment of ROC analysis, and Saito and Rehmsmeier (2015) demonstrate that ROC curves can present an optimistic picture when the positive class

is scarce and that precision-recall information is more diagnostic in such cases. This is why the threshold sweep reported in Table 2 is presented alongside the ROC-AUC rather than in place of it.

Because the strongest predictor in this study turns out to be the volume of Chinese state-media coverage, the literature on how that media system operates is directly relevant. Brady (2008) describes the propaganda apparatus as a coordinated instrument of governance whose output is planned rather than emergent, and Stockmann (2013) shows that although commercialisation has changed the style of Chinese media, the state retains the capacity to direct attention on issues it treats as strategic. King et al. (2017) demonstrate empirically that coordinated state activity in Chinese online media is timed around events of political salience and functions to shape attention rather than to win argument, a finding that supports reading a surge in coverage as a scheduling signal rather than a sentiment signal. Roberts (2018) reaches a compatible conclusion by a different route, showing that information control in China operates as much through the allocation of attention as through outright suppression. Read together, this work supplies the mechanism that the empirical results in Section 4 require: if coverage of strategic minerals is planned and centrally coordinated, then an increase in its volume is a by-product of an internal policy process that is already under way and is therefore available as a leading indicator to an outside observer who is counting articles rather than reading them. Figure 3 sets out the thematic structure of the four strands reviewed above and the position of the present study within them.




Figure 3. Thematic structure of the literature review and positioning of the present study.

Taken as a whole, the literature establishes that Chinese rare earth policy is deliberate and historically patterned, that markets respond to it incompletely, that GDELT offers a usable and public measure of geopolitical activity, and that state-media output is a planned instrument rather than a free market of opinion. What has not been attempted is the combination. No published study, to the knowledge of the author, applies a state-media coverage filter over

GDELT to issue dated, out-of-sample warnings of rare earth export restrictions and then validates those warnings against the announcements that followed. The present study addresses that gap and does so with a deliberately simple model so that the contribution rests on the signal and on the validation design rather than on algorithmic complexity.

Data and Methodology

Data Sources

A total of three open datasets were combined to form a single daily timeline with 1,556 data points for the duration of January 2020 up to December 2025, as stated in Table 1. The information collected through the GDELT v2 Event Database, which was accessed via Google BigQuery, highlighted the daily numbers, average tone, and mean Goldstein conflict score of events sourced specifically from China, with regards to CAMEO event codes related to restrictions (061, 163, 193, 172, 173). The second query was cautious when filtering the URLs be checked to get Chinese state-run media coverage. In this case, the specific URLs (Xinhua, Global Times, MOFCOM, People’s Daily, etc.) had to be filtered to make sure that their information is reliable and relevant. During this query, the numbers (daily totals and tone) of events were gathered separately. Apart from this, Yahoo Finance was used as a source of daily closing prices for REMX, BEL.NS, and DIXON.NS securities. Downloading was made possible through the usage of the yfinance library.

Table 1. Data sources merged into a single daily dataset (2020–2025).

Data source	What it measures	Key fields
GDELT Event Database (Leetaru & Schrod, 2013)	Volume and tone of China-related geopolitical events	Event count, average tone, Goldstein scale
GDELT (state-media filter)	Chinese government media coverage	Xinhua / Global Times / MOFCOM / People’s Daily counts and tone
Yahoo Finance (Aroussi, 2023)	Market reaction	REMX, BEL, Dixon daily close and log returns

Label Definition

A binary label was introduced in which a day was deemed "crisis" (1) if REMX's expected 14-day return exceeds 15% and "normal" (0) otherwise. The 15% threshold was selected to capture true supply-shock responses and not normal market fluctuations. This results in about 130 crisis days compared to 1,400 normal days, creating an imbalance as the crisis takes place 8.5% of the time, which impacts the later selections made in the modeling process.

Feature Engineering

Unrefined counts of daily incidents alone are too unpredictable to have any significance: for instance, on a single busy day of reporting, many stories could have been reported by the same reporter rather than showing an actual change in global events. In order to solve this and catch trends, ten indicators were created, including rolling weekly averages of average events and sentiments, a measure of tone acceleration which shows how fast the attitude changes during a

week, a virus indicator that works only when the daily events volume is over two standard deviations more than usual, and several other indicators showing media activity and changes of the indexes. Due to this, sometimes it is more important to know the tendency of events over a week rather than knowing the volume of them any particular day. Rows, which had some missing cells because of incomplete rolling-window calculations, were removed from this analysis leaving 1,525 days.

Model Training and Evaluation

StandardScaler was utilized to standardize features on training data to effectively avoid the potential leak of information. The logistic regression classifier (`class_weight='balanced'`, `max_iter=1000`, `random_state=42`) was built and trained from the dataset from 2020–2024 (1,266 days, 100 events) and tested with the dataset from 2025 (259 days, 30 events). Logistic regression was selected since the training set included only 130 positive samples, and it is highly possible for more sophisticated algorithms to cause overfitting with such a small amount of training data. The performance was tested on confusion matrix, ROC-AUC, and precision and recall calculation. The operating threshold was chosen to be equal to 0.20 to ensure that recall was more prioritized than precision since the cost of missing an actual restriction event is higher than the cost of a false alarm. Event validation was conducted with the help of looking into the 30-day period before the true event occurrence. The analysis was performed using Python and its libraries such as pandas, scikit-learn (Pedregosa et al., 2011), and matplotlib.

The evaluation design deserves a short justification, because it determines whether the results reported below can be read as genuine forecasts rather than as in-sample description. The split is chronological and single: every observation used for fitting precedes every observation used for testing, and no scaling parameter, feature statistic or threshold was estimated on the 2025 period. Random k-fold cross-validation was deliberately avoided, since shuffling would allow the model to learn from days that follow the days it is asked to predict and would inflate performance in a way that could not be reproduced in deployment (Bergmeir & Benítez, 2012). Class weighting rather than synthetic over-sampling was used to address the 8.5% positive rate, so that no observation entering the model is artificial (Chawla et al., 2002). Because ROC-AUC can be optimistic under strong class imbalance, results are reported at several operating thresholds as well as in aggregate, following the recommendation of Saito and Rehmsmeier (2015). The event-level validation described in Section 4.3 is treated as the stricter of the two tests, since a model may achieve an acceptable aggregate score while still failing to alert before the events that matter.

Results

Classifier Performance

On the unseen 2025 test set the model achieved an ROC-AUC of 0.732 (Figure 4), comfortably exceeding the 0.500 baseline of random guessing. At the chosen threshold of 0.20, recall reached 0.87, identifying 26 of the 30 real crisis days, at the cost of low precision (0.16), generating 138 false positives against 26 true positives (Table 2). The precision-recall trade-off at alternative thresholds is summarised below.




Figure 4. ROC curve on the 2025 hold-out test set. The classifier (AUC = 0.732) separates crisis from normal days better than chance, though it is far from perfect.

Table 2. Threshold sweep on the 2025 hold-out test set. The 0.20 operating point (bold) was chosen to prioritise recall.

Threshold	Recall	Precision	TP	FP	FN
0.50	0.37	0.46	11	13	19
0.40	0.40	0.26	12	34	18
0.30	0.57	0.20	17	69	13
0.20	0.87	0.16	26	138	4
0.15	0.93	0.14	28	170	2

1.1 Feature Importance

The most informative result emerged from the model’s standardised coefficients (Figure 5). The strongest positive predictor of an approaching crisis was the 7-day rolling volume of Chinese state-media coverage, with a standardised coefficient of +1.24. The sheer quantity of government media attention to strategic minerals, rather than the negativity or hostility of that coverage, provided the strongest signal. Sentiment-based features were comparatively weak. One methodological note: the overall GDELT event-count feature carried a large negative coefficient, which most likely reflects collinearity between the two count features rather than a meaningful protective effect; the volume signal, rather than the individual coefficient magnitudes, should be read as the principal finding.




Figure 5. Standardised logistic-regression coefficients. Chinese state-media volume (top) is the strongest positive driver of a crisis alert, while sentiment and price-momentum features matter far less.

Validation on Real Events

The classifier was validated against the two real restriction events of 2025 (Table 3). For the April 2025 restriction of seven rare earth elements, the classifier raised its first alert on 5 March, 30 days before the official announcement, and REMX subsequently rose 11.8% within 21 days. For the October 2025 restriction of twelve elements, the first alert came on 12 September, 27 days before the 9 October announcement, ahead of a 6.6% REMX move. The average lead time across both events was 28.5 days, and in both cases the alerts began while REMX was still flat or declining, before the market had incorporated the information (Figure 6).

Table 3. Classifier alerts ahead of the two 2025 restriction events.

Event	Announceme nt	First alert	Lead time	REMX move (21d)
April 2025 — 7 REEs restricted	4 Apr 2025	5 Mar 2025	30 days	+11.8%
October 2025 —	9 Oct 2025	12 Sep 2025	27 days	+6.6%

Event	Announcement	First alert	Lead time	REMX move (21d)
12 REEs restricted				




Figure 6. Event validation for the two 2025 restrictions. In both panels the crisis probability (blue) crosses the alert threshold (orange) weeks before the official announcement (red), while REMX (black) is still quiet.

Cross-Market Contagion

The 14-day return correlation between REMX and Dixon Technologies nearly tripled, rising from 0.17 on normal days to 0.45 on crisis days (Figure 7). The REMX–BEL correlation rose more modestly from 0.23 to 0.29. This correlation tightening confirms that a rare earth supply restriction acts as a common shock propagating across related equities and across borders into the Indian downstream supply chain (Mancheri et al., 2019), consistent with the cross-sectoral spillover findings of Seiler (2025) and Zhang et al. (2024).




Figure 7. Cross-market return correlation, normal vs. crisis days. Markets that drift apart in calm periods move together during supply shocks, especially REMX and Dixon.

Discussion

The main result is that a versatile and fully reproducible classifier identified two occurrences of rare earth supply restriction in real-time about 28 days before these announcements were made by the authorities, at times when the market had not yet reacted. This confirms the assumption that the military build-up leading to the restriction can be identified utilizing open data and that the period of delay between the Chinese signals and their subsequent impact on the Western markets is suitable for systematic exploitation.

The mechanism that would produce such lead time is administrative rather than mysterious. Export control is the terminal step of a bureaucratic process that involves consultation between ministries, drafting, internal circulation and the preparation of public justification, and the periodisation offered by Shen et al. (2020) suggests that this process has been reasonably consistent in form even as its objectives have changed. Coverage volume rises during that preparatory phase because the apparatus described by Brady (2008) and Stockmann (2013) is being used to establish the framing that the eventual announcement will occupy. An outside observer counting articles is therefore not predicting a decision so much as observing the visible surface of a decision that has already been taken internally but not yet published. This interpretation is consistent with the direction of the coefficients reported in Section 4.2 and, importantly, it is falsifiable: if the signal were noise, it would not concentrate in the weeks immediately preceding announcements, and the event-level validation in Table 3 would not show alerts clustering there. The finding that the media volume aspect of our analysis is more powerful than the media sentiment aspect needs explaining. At first, we expected the model to reveal a negative or hostile form of coverage—negative coverage. Instead, it discovered that the only important indicator is the volume of government attention to strategic minerals, regardless of whether the tone is hostile or not. An increase in the amount of media coverage appears to be an indicator of future policy changes. This fits with the findings by Ferrara et al. (2020) on the superiority of news volume measures over tone measures in GDELT-based forecasting tasks, and the finding of the political science community that authoritarian governments' state-owned media functions only as a signal for upcoming political events (King et al., 2017).

Placing these results next to the existing literature clarifies both what is new and what is confirmatory. The confirmatory element is the cross-sectoral transmission: the correlation tightening reported in Section 4.4 is the same phenomenon that Seiler (2025) identifies through quantile regression and that Zhang et al. (2024) measure as risk spillover, observed here at daily frequency and extended across a border into Indian downstream equities. The new element is temporal. Proelss et al. (2018) established that REE prices break structurally at policy announcements and do not follow a random walk afterwards; the present results suggest that the informational inefficiency they document begins before the announcement rather than at it and can be partially exploited by an observer with access only to public data. Similarly, while Li et al. (2023) demonstrate that political risk explains REE price variation, they do so contemporaneously, whereas the finding here is that a specific and countable component of that risk is observable in advance. The contribution is therefore less a claim to have outperformed

prior models than a claim to have shifted the same relationship from an explanatory to a predictive frame.

There are some limitations that should be mentioned clearly. First, the out-of-sample validation is based on only two real events because the July 2023 gallium and Germanium controls are included in the training period and cannot be used as an unknown test. Two successes, however, are not enough to make conclusions about generalizability. Second, the false-alarm rate is too high. The model creates an alarm for trading every two days if the threshold of 0.20 is used, which, once applied, can lead to unreliability. A real-life implementation presupposes using a higher threshold, another rule of confirmation, or a second model for filtering out false alarms. Third, although REMX is a good medium of the condition of the rare earth market, it is not a significant ratio of the price of the physical commodity. The fourth limitation is that logistic regression is not a complicated method and is not going to discover intensive interaction of variables that would be possible using other methods like Bayesian type ones (Pedregosa et al., 2011). The false-alarm problem deserves a more constructive treatment than a simple acknowledgement, because it determines whether the classifier is usable at all. The threshold sweep in Table 2 makes the trade-off explicit: moving from 0.20 to 0.30 removes half of the false positives but also discards a third of the true events, while moving to 0.15 recovers almost every event at the cost of a near-continuous alarm. Which operating point is correct is not a statistical question but an economic one, and it depends on the loss function of the user. For a procurement manager holding three months of inventory, a false alarm costs the price of an unnecessary forward purchase while a missed event costs a production stoppage, and that asymmetry justifies the recall-weighted setting used here; for a trader carrying the financing cost of a position, the asymmetry runs the other way. A more practical deployment would treat the model output as one input rather than as a decision rule. Requiring the crisis probability to remain above threshold for several consecutive sessions or requiring confirmation from an independent indicator such as inventory drawdown or shipping data, would reduce the alarm rate substantially at a modest cost in lead time. This is the sense in which the precision-recall framing of Saito and Rehmsmeier (2015) is more informative than the headline ROC-AUC: a single aggregate number conceals the choice that has to be made. The importance of an observation on methodology should be underscored, since the factors that caused the results to be accepted or rejected were not based on algorithms, but on methodology of labeling the data defined as well as the accuracy with which the time series data was split into training and test sets. In this study, data was not shuffled so as to prevent the models from using future information and generating impressive yet meaningless results.

Finally, the results carry implications for importing economies that are not visible in aggregate trade statistics. India occupies an intermediate position in this supply chain, importing processed rare earth inputs and finished magnets for downstream electronics and defence manufacturing while holding substantial and largely unexploited reserves of its own. The correlation results indicate that a Chinese restriction is transmitted into Indian downstream equities within days, which means that the exposure is real even though India is not a party to the announcement. An early-warning capability of the type described here would be of most use to precisely such intermediate economies, where the policy responses available in a four-week window - accelerating existing orders, releasing buffer stock, or triggering pre-negotiated

alternative supply arrangements - are meaningful but require advance notice to execute. The same logic applies to the diversification strategies compared by Barteková and Kemp (2016): the instruments differ across importing regions, but every one of them is easier to deploy with lead time than without it. It should be stressed that this is an argument for monitoring rather than for market timing, since the false-alarm rate documented above is far too high to support the latter.

Conclusion and Future Scope

In this research, it has been shown that the classifier made based on the three publicly accessible open datasets - the GDELT global event database, indicators of Chinese state media coverage, and daily stocks prices - can be useful for predicting rare earths supply restriction events. The classifier reached an ROC-AUC of 0.732 on the unseen data of 2025 and provided alerts 28.5 days prior to two actual announcements about restrictions, despite a period of market calmness. As having been determined, the strongest predictive indicator was the amount of coverage which the Chinese media produced rather than the polarity of that coverage, which shows that when there is a significant media attention from the governmental institutions towards strategic minerals, that could trigger the action. Beyond the specific result, the study makes a methodological point about the kind of data that is sufficient for this class of problem. Nothing used here is proprietary: the event data is publicly queryable, the price series are freely downloadable, and the model is a standard linear classifier fitted with close to default settings. The lead time therefore does not derive from privileged information or from algorithmic sophistication, but from the structure of the problem itself, in which an administrative process leaves a countable public trace before it produces a market-moving announcement. This suggests that the more productive direction for further work is not a more elaborate model on the same features, but better features: Chinese-language source text read directly rather than through translation, ministry-level rather than aggregated outlet counts, and physical price series in place of an equity proxy.

The foremost constraint is the limited availability of out-of-sample events for validation purposes. Nonetheless, future endeavors must work on extending the testing set going forward when additional events leading to the restriction happen, providing a more extensive background for claims on generalization. Also, there are a variety of methodological improvements implied: direct scraping of Chinese media rather than relying on English-translated information provided by GDELT; using real price indices of rare earths (for instance, neodymium oxide prices) as dependent variables; checking if the volume-of-coverage signal is valid for other commodities that are crucially related in the strategy such as cobalt, lithium, and copper (Columbia University Center on Global Energy Policy, 2026); and testing of ensembles or sequential models that could help to minimize false alarm rates without sacrificing recall. The framework could also be applied to forecasting policy-related supply shocks in strategic mineral markets that gain more importance (Atlantic Council, 2025; International Energy Agency, 2024; OECD, 2023).

Data and Code Availability

The data used in this research is available to the public. GDELT event data is available through BigQuery (Leetaru & Schrodtt, 2013). Price data has been retrieved from Yahoo Finance through yfinance Python library (Aroussi, 2023). The analysis pipeline is described in Section 3 and organized through a series of reproducible Python scripts.

References

Aroussi, R. (2023). *yfinance: Download market data from Yahoo! Finance API* [Computer software]. Python Package Index. <https://pypi.org/project/yfinance/>

Atlantic Council. (2025). A US framework for assessing risk in critical mineral supply chains [Issue brief]. Atlantic Council.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593–1636. <https://doi.org/10.1093/qje/qjw024>

Barteková, E., & Kemp, R. (2016). National strategies for securing a stable supply of rare earths in different world regions. *Resources Policy*, 49, 153–164. <https://doi.org/10.1016/j.resourpol.2016.05.003>

Baskaran, G., & Schwartz, M. (2025). *The consequences of China's new rare earths export restrictions*. Center for Strategic and International Studies.

Bergmeir, C., & Benítez, J. M. (2012). On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191, 192–213. <https://doi.org/10.1016/j.ins.2011.12.028>

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1), 1–8. <https://doi.org/10.1016/j.jocs.2010.12.007>

Bouoiyour, J., et al. (2026). A foresight study on the geopolitical vulnerabilities of the rare earth supply chain in securing green hydrogen. *Environmental Earth Sciences*, 85.

Brady, A.-M. (2008). *Marketing dictatorship: Propaganda and thought work in contemporary China*. Rowman & Littlefield.

Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225. <https://doi.org/10.1257/aer.20191823>

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357. <https://doi.org/10.1613/jair.953>

China State Council. (2023). Export control regulations on gallium and germanium related items (MOFCOM Announcement No. 23).

China State Council. (2025). *Export control regulations on rare earth elements* (MOFCOM Announcements, April and October 2025).

Columbia University Center on Global Energy Policy. (2026). *Protecting existing US and allied copper smelting capacity* [Policy paper].

Du, S., Zhang, J., Han, Z., Wang, L., & Tang, D. (2024). Measurement and analysis of geopolitical risk in Southeast Asian countries based on GDELT event data. *World Regional Studies*, 33(4), 13–23.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874. <https://doi.org/10.1016/j.patrec.2005.10.010>

Ferrara, E., et al. (2020). Using the GDELT dataset to analyse the Italian sovereign bond market. In *Complex networks and their applications* (Lecture Notes in Computer Science). Springer.

Guo, Y., et al. (2024). Research on the competitive landscape of global rare earth resources demand. *Managerial and Decision Economics*, 45(4).

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95. <https://doi.org/10.1109/MCSE.2007.55>

International Energy Agency. (2024). *The role of critical minerals in clean energy transitions* (World Energy Outlook Special Report). IEA.

King, G., Pan, J., & Roberts, M. E. (2017). How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review*, 111(3), 484–501. <https://doi.org/10.1017/S0003055417000144>

Kuo, J. (2021). Geopolitical risks in the rare earth supply chain. *Strategic Studies Quarterly*, 15(2), 59–83.

Leetaru, K., & Schrodt, P. A. (2013). *GDELT: Global data on events, location and tone, 1979–2012* [Paper presentation]. International Studies Association Annual Convention.

Li, Z. Z., Meng, Q., Zhang, L., Lobont, O. R., & Shen, Y. (2023). How do rare earth prices respond to economic and geopolitical factors? *Resources Policy*, 85, Article 103853. <https://doi.org/10.1016/j.resourpol.2023.103853>

Mancheri, N. A., Sprecher, B., Bailey, G., Ge, J., & Tukker, A. (2019). Effect of Chinese policies on rare earth supply chain resilience. *Resources, Conservation and Recycling*, 142, 101–112. <https://doi.org/10.1016/j.resconrec.2018.11.017>

McKinney, W. (2010). Data structures for statistical computing in Python. In *Proceedings of the 9th Python in Science Conference* (pp. 56–61).

OECD. (2023). *Critical materials outlook*. OECD Publishing.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.

Proelss, J., Schweizer, D., & Seiler, V. (2018). Do announcements of WTO dispute resolution cases matter? Evidence from the rare earth elements market. *Energy Economics*, 73, 1–23. <https://doi.org/10.1016/j.eneco.2018.05.004>

Roberts, M. E. (2018). *Censored: Distraction and diversion inside China's Great Firewall*. Princeton University Press.

Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), Article e0118432. <https://doi.org/10.1371/journal.pone.0118432>

Seiler, M. (2025). Strategic risk spillovers from rare earth markets to critical industrial sectors. *Risks*, 13(3), Article 156.

Shen, Y., Moomy, R., & Eggert, R. G. (2020). China's public policies toward rare earths, 1975–2018. *Mineral Economics*, 33(1), 127–151. <https://doi.org/10.1007/s13563-019-00214-2>

Stockmann, D. (2013). *Media commercialization and authoritarian rule in China*. Cambridge University Press.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3), 1139–1168. <https://doi.org/10.1111/j.1540-6261.2007.01232.x>

Vekasi, K. (2019). Politics, markets, and rare commodities: Responses to Chinese rare earth policy. *Japanese Journal of Political Science*, 20(1), 2–20. <https://doi.org/10.1017/S1468109918000348>

Wang, Y., et al. (2025). Geopolitical risks and firm export product quality. *International Review of Financial Analysis*. Elsevier.

Ward, M. D., Greenhill, B. D., & Bakke, K. M. (2010). The perils of policy by p-value: Predicting civil conflicts. *Journal of Peace Research*, 47(4), 363–375. <https://doi.org/10.1177/0022343309356491>

Zhang, X., et al. (2024). The risk spillover between geopolitical risk and China's 5G, semiconductor and rare earth industries. *Heliyon*, 10(23).

Zhou, B., Li, Z., & Chen, C. (2017). Global potential of rare earth resources and rare earth demand from clean technologies. *Natural Resources Research*, 27, 201–216. <https://doi.org/10.1007/s11053-017-9346-3>