AI Sovereignty

Zhenyu Gao[1], Wenxi Jiang[1], Haohan Ren[2], Kemin Wang[3], and Yuezhi Wu[3]

[1]Chinese University of Hong Kong

[2]Zhejiang University of Finance and Economics

[3]Fudan University

CESI

August 14, 2026

Technology Adoption is Political

> *"There is currently very little work on the political economy of technology choices"*
>
> *"A particularly interesting direction would be to investigate how political economy factors influence the direction of technology..."*
>
> *Acemoglu (2025 AER)*

Technology adoption is not only about efficiency; it can also be shaped by political incentives and state power.

Motivation

GenAI is a Sharp Setting

- **Productivity gains:** a general-purpose technology with the potential to reshape production, management, and innovation.
- **Geopolitical competition:** governments increasingly care not only whether firms adopt AI, but which AI ecosystem they adopt into.

- **AI sovereignty**: states foster and control domestic AI ecosystems to preserve autonomy, security, and strategic capacity.

To study the political economy of GenAI choices,
China's **lag-then-catch-up episode** provides a powerful setting.

Lag Episode and Sovereignty Pressure

Q1 2023: ChatGPT API

- Expanded commercial access to a leading foreign model.
- Firms could access it through third-party cloud services.
- Authorities emphasized cross-border data leakage and foreign system dependence risks.

慎用ChatGPT！中国支付清算协会：智能化工具已暴露跨境数据泄露风险 原创

红星新闻 2023-04-11 14:35

ChatGPT最近很火，而它背后潜藏的风险也引起了支付行业的特别关注。红星新闻记者11日获悉，中国支付清算协会已在其官网发出倡议，希望大家谨慎使用ChatGPT等工具。

中国支付清算协会表示，近期，ChatGPT等工具引起各方广泛关注，已有部分企业员工使用ChatGPT等工具开展工作。但是，此类智能化工具已暴露出跨境数据泄露等风险。为有效应对风险、保护客户隐私、维护数据安全，提升支付清算行业的数据安全管理水平，中国支付清算协会向行业发出倡议——

支付行业从业人员要遵守所在地区的网络使用规定，正确认识ChatGPT等工具的机遇和风险，全面评估使用ChatGPT等工具处理工作内容的风险，依法合规使用ChatGPT等工具。

Catch-Up and State Promotion

Q1 2023: ChatGPT API

- Leading foreign model with large business potential.
- Access existed, but data and control concerns were salient.

慎用ChatGPT！中国支付清算协会：智能化工具已暴露跨境数据泄露风险 原创

红星新闻 2023-04-11 14:35

ChatGPT最近很火，而它背后潜藏的风险也引起了支付行业的特别关注。红星新闻记者11日获悉，中国支付清算协会已在其官网发出倡议，希望大家谨慎使用ChatGPT等工具。

中国支付清算协会表示，近期，ChatGPT等工具引起各方广泛关注，已有部分企业员工使用ChatGPT等工具开展工作。但是，此类智能化工具已暴露出跨境数据泄露等风险。为有效应对风险、保护客户隐私、维护数据安全，提升支付清算行业的数据安全管理水平，中国支付清算协会向行业发出倡议——

支付行业从业人员要遵守所在地区的网络使用规定，正确认识ChatGPT等工具的机遇和风险，全面评估使用ChatGPT等工具处理工作内容的风险，依法合规使用ChatGPT等工具。

Q1 2025: DeepSeek

- Comparably strong domestic LLM; first to rank #1 on major LLM leaderboards.
- State promotion became salient after Xi Jinping met Liang Wenfeng.
- Sovereignty-friendly adoption environment.

From Courtrooms to Crisis Lines, Chinese Officials Embrace DeepSeek

Government bodies nationwide have been eager to show they are using DeepSeek's A.I. technology since the company's founder met with Xi Jinping, China's leader.

This Paper: From Contrast to Test

The ChatGPT–DeepSeek contrast lets us test whether sovereignty exposure changes firms' technology choices.

- **Null hypothesis:** if firms care only about productivity, sovereignty exposure should not predict differential adoption around foreign versus domestic model shocks.

1. Do sovereignty-exposed firms delay GenAI adoption after ChatGPT and accelerate adoption after DeepSeek?
2. What mechanisms explain the delay-and-catch-up pattern?
3. How do investors price sovereignty-(mis-)aligned AI adoption, and why?

How We Answer the Question

We use complementary adoption measures around the two model-release episodes.

- **Investor Interaction Platforms:** extract model-specific adoption from investor-firm dialogues.
  - High-frequency
  - Model-specific
  - Firm-verified responses under exchange oversight
- **GenAI job postings:** use labor-demand evidence as an external revealed-action measure.
- **Cross-sectional differences:** compare high- and low-sovereignty firms around foreign and domestic model shocks.

Main Findings

Sovereignty exposure changes firms' GenAI adoption timeliness, which models they adopt, and how markets price the choice.

- **Delay and catch-up:** sovereignty-exposed firms lag after ChatGPT, then catch up after DeepSeek.
- **Model direction:** the delay is specific to foreign GPT adoption, while the catch-up is driven by DeepSeek.
- **Mechanism:** data-security-exposed firms show stronger delay-and-catch-up patterns; high-sovereignty GPT adopters also avoid data-sensitive uses.
- **Consequences:** investors discount GPT adoption by high-sovereignty firms, as confirmed by their weaker subsequent policy support and performance outcomes.

Related Literature

- **Political economy of technology:** Acemoglu (2025), Beraja, Yang, and Yuchtman (2023, 2024), Chang, Cong, Wang, and Zhang (2025).
- **Geoeconomic conflict and decoupling:** Han, Jiang, and Mei (2024), Clayton et al. (2025), Crosignani et al. (2024).
- **Data and digital sovereignty:** Farboodi and Veldkamp (2021), Chang et al. (2025).
- **AI adoption and firm outcomes:** Bick et al. (2024), Bonney et al. (2024), Fradkin (2025), Babina et al. (2024), Eisfeldt et al. (2023).
- **LLMs in financial text:** Jha, Qian, and Weber (2024), Sheng, Sun, and Zhang (2024).

The contribution is to put political economy into a major technology adoption decision and trace its consequences.

Part I: Measuring adoption

Complementary evidence from investor-firm dialogues and GenAI job postings

Data: Investor Interaction Platforms (SHE and SZE)

- **Official:** responses are mandatory once investors ask, and are subject to securities disclosure rules and exchange oversight.
- **Model-specific:** investors ask whether firms use specific models such as ChatGPT or DeepSeek.
- **High-frequency:** required to respond within business days.
- **Coverage:** 28,822 GenAI-related dialogues across 3,656 firms, covering nearly 70% of listed firms and about 90% of market capitalization.

IIP gives us timely, model-specific firm disclosures, but it is not a stand-alone census of adoption.

Examples: Investor Questions and Firm Responses

Firm	Investor question	Firm response
Guoguang Elec. 002045	It is said that the company, together with international enterprises, has developed the world's first AI smart speaker based on **ChatGPT**, directly integrating **ChatGPT**. What is the progress of this project? When will it be released?	The company ... has already released a video showcasing the **ChatGPT**-enabled smart speaker developed, designed, and produced by the company.
Transfar Log. 002010	... Has it integrated **DeepSeek**?	... Regarding the **DeepSeek** model, the company has recently completed localized deployment, currently applied in campus customer service and dedicated line intelligent quotation scenarios...
China Great Wall 000066	Does the company have any plans in the field of OpenAI's **ChatGPT**? Has **ChatGPT** been integrated? Thank you.	... The company has not yet integrated **ChatGPT**...

Variable Construction: IIP

The measurement pipeline turns unstructured Q&A into a firm-month panel.

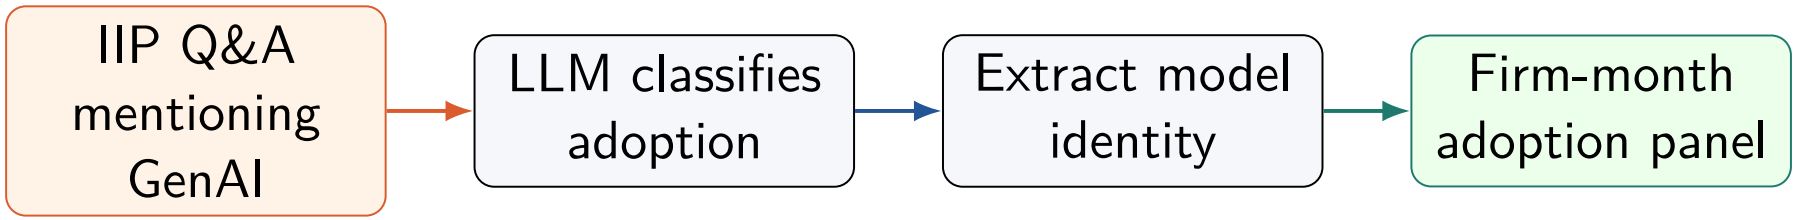

- $AI_Adoption_{i,t} = 1$ if firm i has adopted GenAI by month t.
- Model-specific flags like $GPT_{i,t}$ and $DeepSeek_{i,t}$.
- Adoption is treated as persistent once disclosed.

Next we validate our measure.

Broad and Detailed Application Domain

If our IIP-based adoption captures real activities, then firms should describe concrete applications across business-relevant domains, rather than generic chatbot use that may simply serve investor-relations purposes.

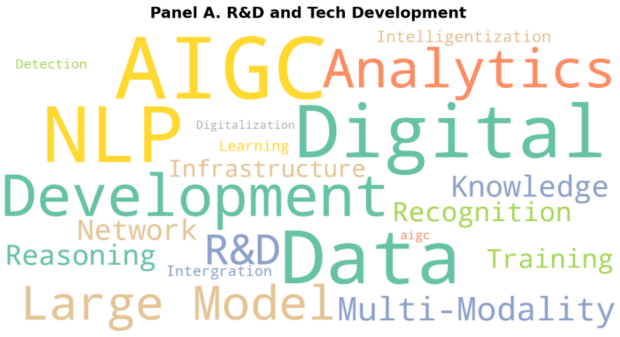

R&D and technology

Operations and management

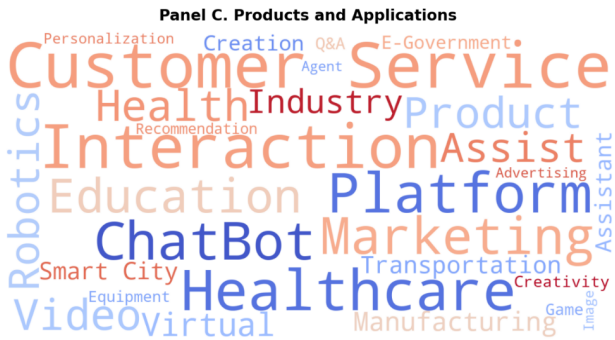

Products and applications

The disclosures point to specific, business-relevant deployment across R&D, operations, and products.

Against Other Real Activities

If the measure captures real activities, then it should be associated with other revealed AI activities.

	(1)	(2)	(3)	(4)	(5)	(6)
	Log(1+#AI)		$I(Algorithm > 0)$		$I(AIJobPosting > 0)$	
AI_Adoption	**1.654*** **	**0.169*** **	**0.032*** **	**0.029*** **	**0.178*** **	**0.094*** **
	(29.76)	**(8.79)**	**(6.08)**	**(6.12)**	**(10.88)**	**(7.51)**
$Log(1+\#AI)_{t-1}$		0.888***				
		(192.41)				
$I(Algorithm > 0)_{t-1}$				0.633***		
				(3.30)		
$I(AIJobPosting > 0)_{t-1}$						0.440***
						(43.26)
Observations	10,633	10,222	10,977	10,977	10,977	10,977
R-squared	0.149	0.811	0.028	0.076	0.014	0.190
FE			Year			

IIP disclosures line up with other AI-related actions, but we still treat job postings as a complementary adoption source.

Data: GenAI Job Postings

We complement IIP disclosures with a labor-demand measure from listed-firm recruitment postings.

- **Source:** all available postings collected from major online recruitment platforms, including Liepin, 51Job, Zhaopin, Boss Zhipin, and Lagou.
- **Text fields:** job title and job description; each posting is assigned to a firm-month.

	IIP	**Job Posting**
A-share coverage	3,656	3,318
GPT adopter	410	185
DeepSeek adopter	855	180

For general adoption, we use both IIP and job postings; for model-specific information, we mainly rely on IIP.

Variable Construction: GenAI Job Postings

We adapt Sautner et al. (2019JF)'s keyword-discovery logic to identify GenAI labor demand in job ads.

- **Seed words:** start from explicit GenAI anchors such as AIGC, large models, LLM, ChatGPT/GPT-4, text-to-image, and text-to-video.
- **Discovered keywords:** expand the vocabulary using relative-rate discovery; examples include multimodal, Prompt, prompt engineering, RAG, LoRA, LangChain, model inference, AI agent, and vector search.
- **Relative-rate discovery:** extract candidate phrases and rank them by

$$\log\left(\Pr(\text{phrase} \mid \text{seed-hit}) / \Pr(\text{phrase} \mid \text{background})\right).$$

- **Final variable:** count GenAI keyword hits in job posting and aggregate to firm-month *GenAIJobPosting* and count of model-specific hits like *DeepSeekJobPosting* and *GPTJobPosting* .

Job postings measure general GenAI hiring demand; model-specific adoption is measured mainly from IIP.

Variable Construction: Sovereignty Exposure

We draw firm-level sovereignty exposure from two literatures: political economy of state influence and U.S.-China Tension.

Political Economy

- **SOE:** direct state ownership and policy alignment (Bai, Lu, and Tao, 2006).
- **Regulated industry:** strategic sectors under stronger policy control, such as energy, telecom, and defense (State Council, 2006; Bradshaw et al., 2019).

U.S.-China Exposures

- **High decoupling:** lower reliance on U.S. patent knowledge, following Han, Jiang, and Mei (2024).
- **Low UCT sensitivity:** low stock-return sensitivity to U.S.-China tension (Rogers et al., 2024), interpreted as detachment from U.S.-related risks.

$$SovIndex = 0.663Z(SOE) + 0.648Z(RegInd) + 0.306Z(HighDecouple) + 0.215Z(LowUCTSensitivity).$$

$HighSovIndex = 1$ if $SovIndex$ is above the sample median.

Part II: Main adoption evidence

Foreign delay, domestic catch-up

Empirical Design

Compares cross-sectional differences around the two model-release shocks.

$$Y_{i,t} = \alpha + \beta\, SovProxy_i \times PostAI_t + \mu_i + \delta_t + \varepsilon_{i,t}.$$

- **IIP-based AI adoption:** cumulative adoption indicator from official investor-firm Q&A, observed at high frequency and by model name.
- **GenAI job posting:** firm-month count of GenAI-related hiring demand from major recruitment platforms.
- **Shock windows:** GPT API release and DeepSeek release are estimated separately.
- **Preferred controls:** firm FE plus time-industry and time-city FE.

If sovereignty matters,
$\beta < 0$ after ChatGPT and $\beta > 0$ after DeepSeek.

IIP Visibility and Sample Strategy

IIP adoption is observed through investor questioning and firm disclosure, so visibility selection is part of the empirical design.

- **Asked by investors Q:** main comparisons are within the questioned sample and tight event windows.
- **Firm response R:** about 97% of firms respond once asked; exchange platforms require timely responses.
- **Selection into being asked:** the GPT and DeepSeek event samples have similar distributions of sovereignty and firm characteristics.
- **Complementary revealed-action measure:** replicate the pattern using GenAI job postings.

IIP is not a population-wide adoption census, but a pure visibility-selection explanation is unlikely to drive the full pattern.

Sample Construction

To reduce concerns that the results are driven by differences in investor attention, we use tight event windows and compare firms within the questioned sample.

- Two tight event windows: $[-6, +6]$ months around February 2023 for GPT and February 2025 for DeepSeek.
- **Within questioned sample:** compare firms that investors asked about GenAI within the relevant window.
  - **GPT window:** firms asked by investors during August 2022 to August 2023.
  - **DeepSeek window:** firms asked both before and after February 2025.

The design keeps the comparison inside firms already attracting investor attention to GenAI.

Similar Distribution across Two Samples

The GPT and DeepSeek event samples look similar on key sovereignty and firm characteristics.

	GPT sample	**DeepSeek sample**
SOE	23.44%	26.28%
RegInd	5.74%	5.96%
HighDecouple	50.07%	48.53%
LowUCTSensitivity	56.19%	51.70%
HighSovIndex	27.50%	28.84%
$\log(AT)$	22.31	22.43
MB	3.91	3.31
Age	21.89	23.57

The two event samples look similar on key sovereignty and firm characteristics.

Gap Widening across Sovereignty Proxies

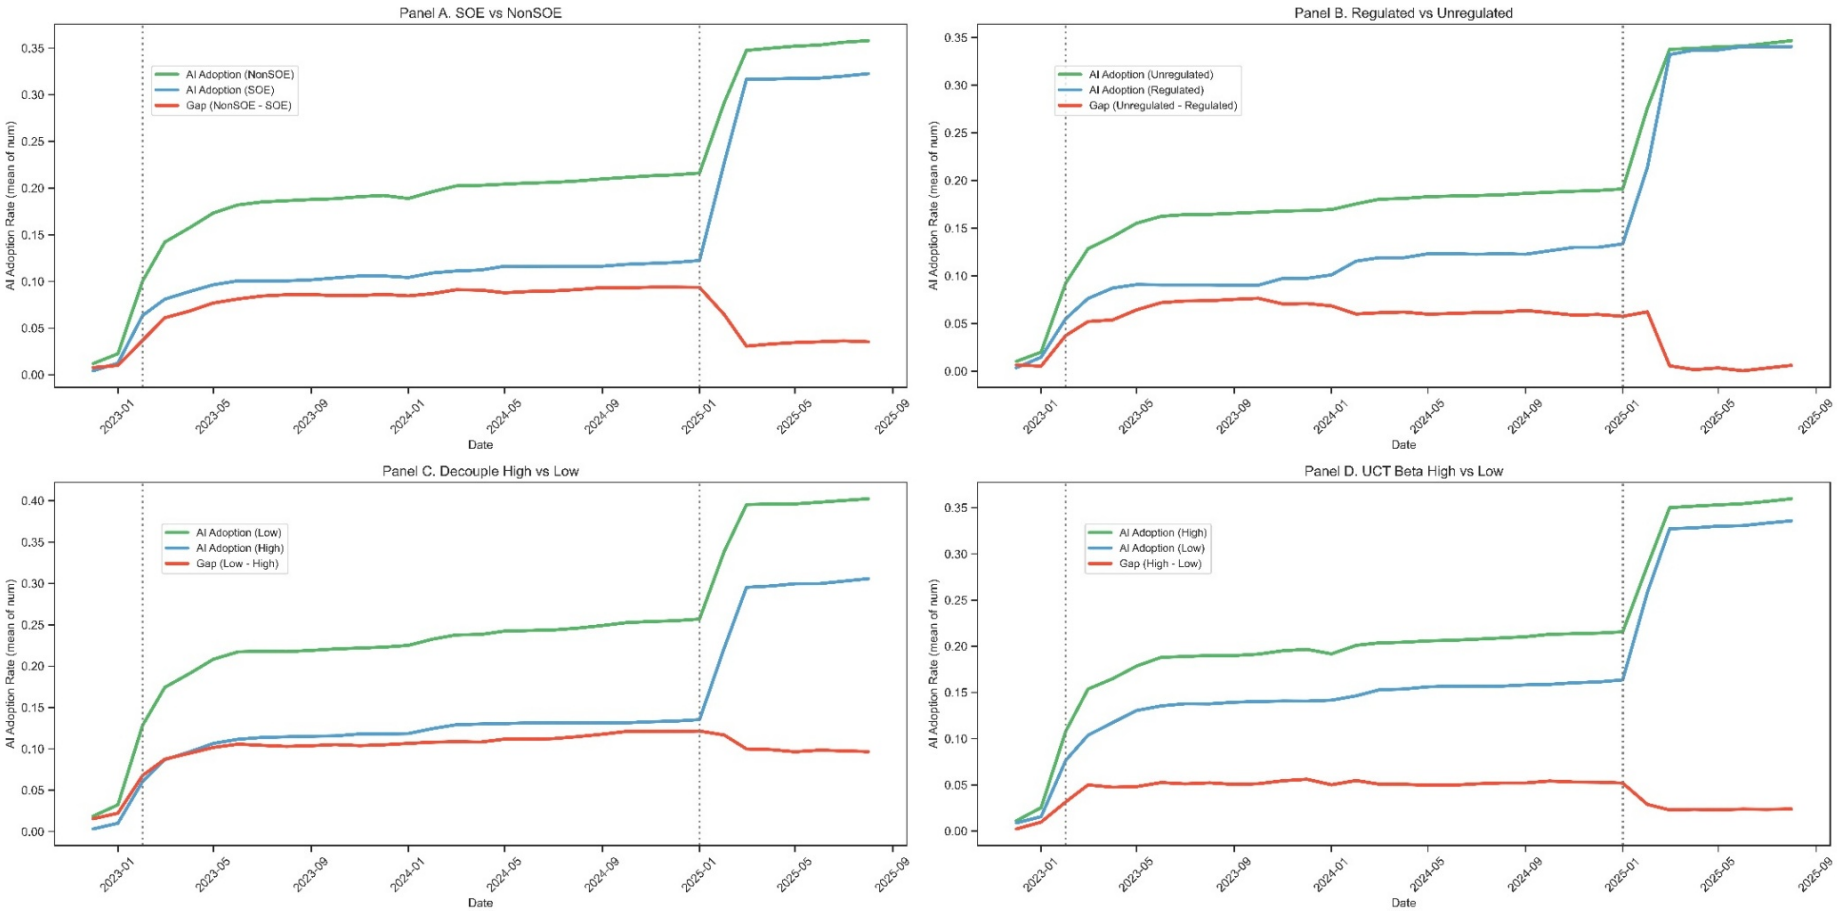

Main Results: Regressions

The regression evidence formalizes the visual pattern

Y= *AI_Adoption*

Panel A: ChatGPT shock				
Sovereignty proxy	SOE	Regulated	Decoupled	Low UCT sens.
SovProxy × *PostGPT*	**-0.058****	**-0.025**	**-0.089*****	**-0.056*****
	(-2.43)	**(-0.57)**	**(-4.41)**	**(-2.72)**
FE			Firm, YearMonth	
Observations	19,012	21,037	19,474	19,474
R-squared	0.627	0.595	0.622	0.620
Panel B: DeepSeek shock				
Sovereignty proxy	SOE	Regulated	Decoupled	Low UCT sens.
SovProxy × *PostDS*	**0.086*****	**0.076***	**0.032***	**0.044****
	(3.86)	**(1.76)**	**(1.78)**	**(2.49)**
FE			Firm, YearMonth	
Observations	19,379	19,542	19,542	19,542
R-squared	0.858	0.856	0.856	0.856

Unified Sovereignty Index

If the four proxies capture a common sovereignty exposure, the PCA index should reproduce the same sign reversal.

	ChatGPT shock		DeepSeek shock	
HighSovIndex × *PostAI*	**-0.052****	**-0.062****	**0.085*****	**0.063****
	(-2.29)	**(-2.19)**	**(3.97)**	**(2.38)**
Observations	19,542	19,366	21,037	19,463
R-squared	0.857	0.872	0.596	0.685
FE	Firm; Time	Firm; Time-Ind; Time-City	Firm; Time	Firm; Time-Ind; Time-City

The sign reversal appears with standard firm-time fixed effects and survives finer controls for local and industry-specific shocks.

Dynamic Evidence: No Pre-Trend, Then Divergence

If model releases drive the gap, high-sovereignty firms should diverge after the shocks, not before them.

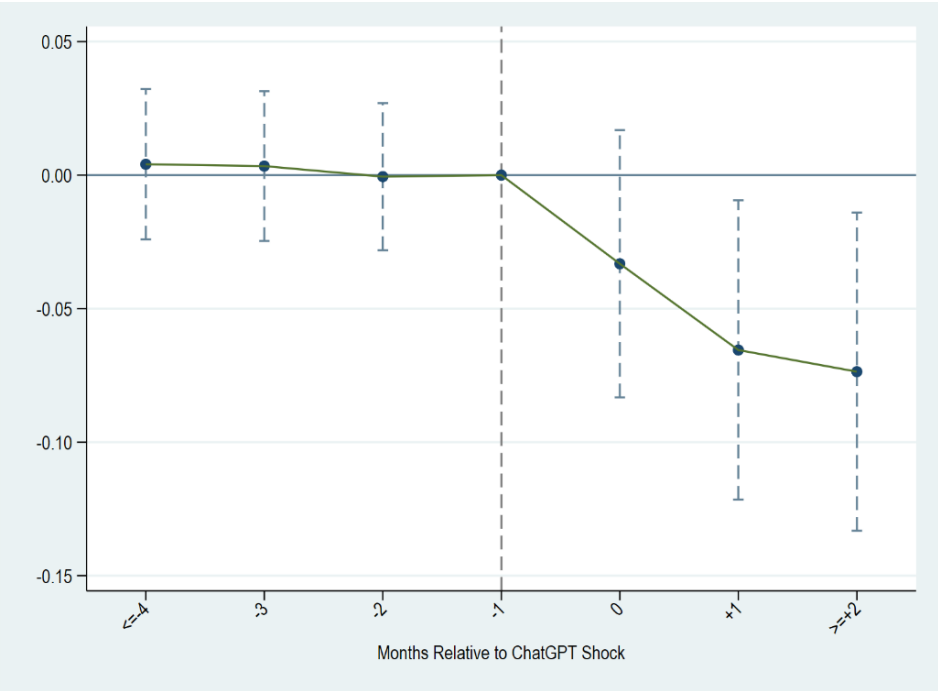

Relative to ChatGPT shock

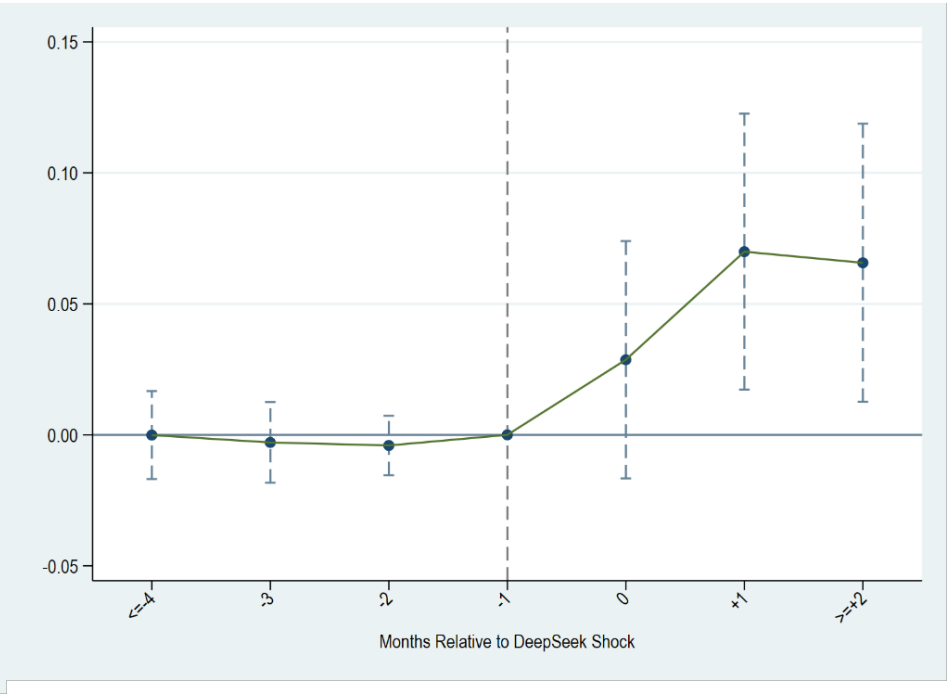

Relative to DeepSeek shock

High-sovereignty firms diverge after the release, not before it.

Model-Specific Adoption Sharpens the Claim

If sovereignty changes model choice rather than generic AI adoption, high-sovereignty firms should move away from GPT and toward DeepSeek.

	Adopt GPT	Adopt DeepSeek
HighSovIndex × *PostAI*	-0.050**	0.062*
	(-2.00)	(1.67)
FE	Firm; Time-Ind; Time-City	
Observations	19,366	11,174
R-squared	0.634	0.644

This moves the interpretation from "AI adoption timeliness" to "direction of model choices"

Complementary Evidence: GenAI Job Postings

A labor-demand measure should show the same lag-then-catch-up pattern if the result is not only IIP visibility selection.

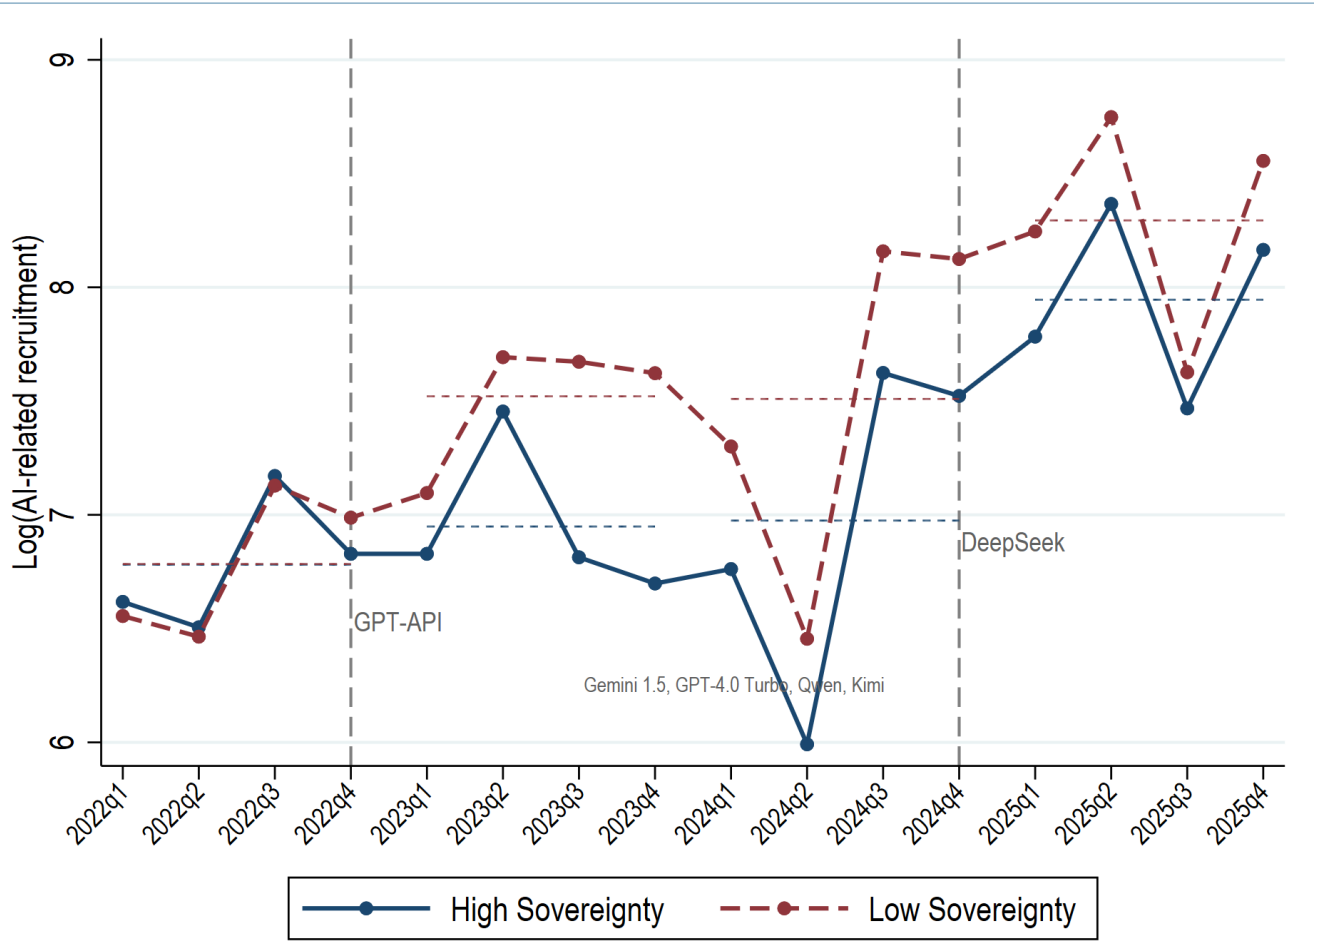

Part III: Why Sovereignty Matters

Data Security and Deployment Risk

Data Security: IIP Evidence

If data security is the relevant wedge, the delay-and-catch-up pattern should be stronger among firms with greater data-security exposure.

Panel A: Keyword-Based Data-Security Exposure				
	ChatGPT shock		DeepSeek shock	
	Low data security	High data security	Low data security	High data security
SovIndex × *PostAI*	-0.070*	-0.098**	0.039	0.093**
	(-1.66)	(-2.14)	(1.05)	(2.32)
High–Low diff.	**-0.028***		**0.054*****	
Diff. P-value	**p < .10**		**p < .01**	
Panel B: LLM-Based Data-Security Score				
	ChatGPT shock		DeepSeek shock	
	Low data security	High data security	Low data security	High data security
SovIndex × *PostAI*	-0.048	-0.088**	0.029	0.064*
	(-1.05)	(-2.22)	(0.84)	(1.95)
High–Low diff.	**-0.040****		**0.035****	
Diff. P-value	**p < .05**		**p < .05**	
FE	Firm; Time-Ind; Time-City			

IIP-based adoption shows stronger GPT delay and DeepSeek catch-up among data-security-exposed firms.

Data Security: Job-Posting Evidence

The same data-security split should appear in revealed labor demand, not only in IIP disclosures.

- We split firms by two data-security proxies: keyword-based exposure and an LLM-based data-security score.

	Keyword-based data-security exposure				LLM-based data-security score			
	ChatGPT shock		DeepSeek shock		ChatGPT shock		DeepSeek shock	
	Low	High	Low	High	Low	High	Low	High
SovIndex × *PostAI*	-0.195	-0.279***	0.054	0.122	-0.173	-0.272**	0.095	0.167**
	(-1.53)	(-2.58)	(0.69)	(1.58)	(-1.55)	(-2.29)	(1.18)	(2.14)
High–Low diff.	**-0.084****		**0.068****		**-0.099****		**0.072****	
Diff. P-value	**$p < .05$**		**$p < .05$**		**$p < .05$**		**$p < .05$**	
Observations	12,898	12,391	28,970	24,736	18,824	7,440	40,056	14,903
Pseudo R^2	0.281	0.277	0.346	0.361	0.276	0.269	0.337	0.371
FE	Firm; YearMonth-Ind; YearMonth-City							

Job postings replicate the data-security pattern using both keyword and LLM-based splits.

Data Security: Intensive Margin

If foreign AI creates data-security risks, high-sovereignty GPT adopters should avoid data-sensitive applications even conditional on adoption.

Y= application data-risk score Higher values = more data-sensitive use		
	GPT adopters	DeepSeek adopters
HighSovIndex	**-0.301*** **	0.031
	(-4.27)	(0.33)
Observations	405	741
R-squared	0.018	0.000

- High-sovereignty GPT adopters use GPT in lower-risk domains.
- No comparable risk shift appears for DeepSeek adopters.

Foreign AI is allocated away from data-sensitive uses; domestic AI is not.

Part IV: Market reaction and real consequences

How investors price sovereignty-sensitive adoption, and why

Market Reaction: Adoption Premium

Benchmark: low-sov firms adoption reflects valuation of model adoption, with limited sovereignty pressure.

If sovereignty concerns matter, high-sovereignty firms should receive a lower premium for GPT adoption.

	Low sovereignty	High sovereignty	Premium gap
ChatGPT adoption			
Adopt GPT	1.545%	0.236%	
Not adopt GPT	0.551%	0.333%	
Adoption premium	0.994%***	-0.096%	**1.090%***
P-value	(0.001)	(0.437)	**(0.066)**
DeepSeek adoption			
Adopt DeepSeek	0.369%	0.447%	
Not adopt DeepSeek	-0.278%	0.202%	
Adoption premium	0.646%***	0.245%	0.401%
P-value	(0.007)	(0.198)	(0.167)

Relative to the low-sovereignty benchmark, high-sovereignty firms are penalized for adopting GPT, but not for adopting DeepSeek.

Market Reaction within Adopters

DeepSeek provides a sovereignty-friendly benchmark among adopters.
A larger high-low gap for GPT adoption identifies the market penalty from sovereignty mismatch.

	DeepSeek adoption	GPT adoption	Gap difference
Low-sovereignty firms	0.368%	1.545%	
High-sovereignty firms	0.447%	0.236%	
Low–high gap	-0.079%	**1.309%****	**1.388%****
P-value	(0.420)	**(0.034)**	**(0.033)**

- The sovereignty gap is large and significant among GPT adopters.
- No comparable gap appears among DeepSeek adopters.

The market discount is tied to sovereignty mismatch, not to GenAI adoption per se.

Regression Evidence

Panel A: Adopt vs. non-adopt sample

Y= CAR[-1,1], FF5		
	GPT	DeepSeek
HighSov × *Adopt*	**-0.015****	-0.003
	(-2.12)	(-0.88)
Adopt	0.009***	0.005**
	(2.58)	(2.00)
HighSov	-0.001	0.005**
	(-0.19)	(2.27)
Observations	1,626	2,572
R-squared	0.119	0.100
FE	YM, Weekday, Industry, City	

Panel B: Within adopters

Y= CAR[-1,1], FF5	
HighSov × *GPT_Over_DeepSeek*	**-0.023*** **
	(-2.79)
GPT_Over_DeepSeek	0.034**
	(2.29)
HighSov	0.005
	(1.15)
Observations	1,140
R-squared	0.154
FE	YM, Weekday, Industry, City

- Results robust to abnormal returns measured by CAPM, FF3 or CN4 (Liu et al., 2019).

High-sovereignty firms receive a GPT discount after controlling for YearMonth, Weekday, Industry, and City FE.

Real Consequences: Weaker Policy Support

When AI adoption is misaligned with the state's sovereignty strategy, high-sovereignty firms may receive weaker policy support.

Prior work motivates two policy outcomes: certification-based industrial support, captured by SRDI or HTE status (Chen et al., 2021; Cao et al., 2024), and discretionary fiscal treatment, captured by tax refunds (Wu et al., 2012; Lin et al., 2018).

	GPT adopters		DeepSeek adopters	
	SRDI or HTE	Change in Tax Refund/AT	SRDI or HTE	Change in Tax Refund/AT
HighSov × *Adopt*	**-0.076***	**-0.001****	-0.034	0.000
	(-1.77)	**(-2.36)**	(-1.12)	(0.34)
Adopt	**0.051***	**0.000****	0.019	-0.000
	(1.73)	**(2.33)**	(0.75)	(-1.39)
Comparison within adopters				
HighSov × *GPT_Over_DeepSeek*	**-0.088****	**-0.001****		
	(-1.98)	**(-2.56)**		
GPT_Over_DeepSeek	-0.006	**0.002****		
	(-0.08)	**(2.56)**		
FE	Industry, city, time			

Sovereignty-misaligned GPT adoption is followed by weaker state-linked policy

Real Consequences: Weaker Performance

If weaker policy support translates into real consequences, high-sovereignty GPT adopters should see weaker performance.

	GPT adopters			DeepSeek adopters		
	Change in Analyst Rating	Change in Forecast EPS	Change in Revenue/AT	Change in Analyst Rating	Change in Forecast EPS	Change in Revenue/AT
HighSov × *Adopt*	**-0.124***	**-0.081***	**-0.090*****	-0.011	-0.022	0.012
	(-1.74)	**(-1.78)**	**(-3.16)**	(-0.28)	(-0.63)	(0.63)
Adopt	0.012	0.003	0.025*	0.027	-0.011	-0.019
	(0.33)	(0.10)	(1.67)	(0.92)	(-0.42)	(-1.34)
Comparison within adopters						
HighSov × *GPT_Over_DeepSeek*	**-0.151****	**-0.095***	**-0.040***			
	(-2.00)	**(-1.72)**	**(-1.71)**			
GPT_Over_DeepSeek	-0.061	0.194*	0.006			
	(-0.89)	(1.77)	(0.33)			
FE	Industry, city, time					

The market aligns with the state,
in anticipation of policy and performance costs.

Summary

- The paper provides high-frequency evidence that state power shapes corporate GenAI adoption.
- Sovereignty-exposed firms delay foreign AI adoption after ChatGPT and catch up after DeepSeek.
- Investors discount GPT adoption by high-sovereignty firms, consistent with policy and performance penalties.

Sovereignty pressure turns GenAI adoption into a political-economy equilibrium.

Implications

The findings matter beyond Chinese firms' GenAI adoption.

- **Fragmented tech diffusion and divergent data-model loops:** frontier technology may diffuse along national blocs rather than purely by technical fit; as firms sort into different model ecosystems, data flows and future model improvement can also diverge.
- **Firms as sovereignty agents and market alignment:** firms carry national AI strategy into technology choices, while investors help discipline misaligned adoption when it predicts weaker policy support and performance.