



OpenClaw and the "Little Lobster" Fever: A Qualitative Study of China's 2026 AI-Agent Adoption Wave

Lihui Zhang*

M-RCBG Senior Fellow at Harvard Kennedy School, Former President of Caixin Media

* Corresponding Author Email: lihu2i@gmail.com- ORCID: 0000-0002-5247-7436

Article Info:

DOI: 10.22399/ijcesn.5500

Received : 20 July 2026

Revised : 02 September 2026

Accepted : 05 September 2026

Keywords

AI agents,
agentic AI,
platform governance,
digital labor,
technology adoption,
China

Abstract:

In March 2026, thousands of people queued outside technology campuses in Shenzhen, Beijing, and Hangzhou to have an open-source agent framework called OpenClaw installed on their machines. Chinese users renamed the tool the "Little Lobster," and the act of running it became "lobster farming." This paper treats that episode as a case worth explaining rather than a curiosity to report. Using qualitative document analysis of policy texts, corporate announcements, community discussion, and reported adopter accounts, it asks who joined the wave, why the Chinese deployment of agent software diverged so sharply from its Western counterpart, and what the episode reveals about how a large economy governs a fast-moving technology. The analysis finds a single social driver beneath four visible tensions, a policy structure that pairs central-level red lines with local-level competition to attract investment, and a corporate value cycle that turns individual anxiety into platform revenue. I introduce three interpretive concepts, radical pragmatism, algorithmic survivalism, and digital narcotization, to name the mechanisms at work, and I argue that the phenomenon is best read as a stress test that China is running on behalf of the wider world, at a cost borne mainly by ordinary users. The paper closes by sketching the likely path from an improvised, command-line era toward regulated, platform-native agent services.

1. Introduction

A queue is a simple thing, and usually a boring one. The queues that formed on March 6, 2026 outside Tencent's Shenzhen headquarters, Baidu's Beijing campus, and the Alibaba Cloud Valley center in Hangzhou were not boring. Close to a thousand people waited at each site, not for a new phone or a pair of limited sneakers, but for a volunteer engineer to install a piece of open-source software on a laptop. The software was an AI agent framework. Within days the queues, the software, and the people in them had become a national talking point.

The framework was OpenClaw, released a few months earlier under the name Moltbot by the Austrian developer Peter Steinberger. Its logo shows a pair of red lobster claws, and Chinese users seized on the image, calling the tool the "Little Lobster" and describing the work of setting it up as "lobster farming." The greeting "have you farmed your lobster today" spread across social feeds as a marker of belonging. What people were installing

was not a chatbot. An agent of this kind takes control of a computer and carries out tasks: it opens a browser to look things up, writes and runs local scripts, sorts files, answers email, and connects to the applications, WeChat and Feishu among them, where Chinese business actually happens. The difference from an ordinary large language model is the difference between advice and action. A chatbot tells you what to do. An agent does it.

This paper takes the episode seriously as an object of study. The queues are easy to describe and hard to explain, and the explanation matters because the same agent software spread very differently in China than it did elsewhere. The goal here is not to celebrate the wave or to warn against it, but to account for its shape: who joined it, what made the Chinese version distinct, how the state responded, and what the whole affair says about governing a technology that moves faster than the rules written for it.

2. Background: what OpenClaw

OpenClaw belongs to a class of systems that researchers call autonomous or multi-agent software, where a program is given a goal and the latitude to pursue it through a sequence of its own actions rather than a single response [1]. The recent version of the idea rides on large language models, which supply the reasoning layer, and adds tool use, memory, and permission to touch the file system, the network, and other applications. That last capability is what makes an agent useful and what makes it dangerous. A system that can read your inbox and run commands on your behalf can also be steered by text it reads, a failure mode known as indirect prompt injection, in which instructions hidden in a web page or a document hijack the agent's behavior [2]. OpenClaw shipped with broad default permissions, which lowered the barrier to doing real work and raised the exposure to exactly this kind of attack.

The name change from Moltbot to OpenClaw was more than branding. The lobster image gave the tool a face, and the face gave a difficult idea a handle that non-technical people could hold. Renaming a hard concept in familiar terms is a well-documented way that technical ideas move through a population [3]. When a shop owner tells a friend that he is farming lobsters rather than deploying an agent, the fear attached to the technology drains out and something closer to a game takes its place. The meme did real work in lowering the cost of adoption.

3. Aim and research questions

The study pursues three questions. The first is descriptive: which groups drove the adoption wave, and what motivated them? The second is comparative: why did the Chinese deployment of agent software follow a different path from the Western one, and what does that difference consist of? The third is interpretive: what does the episode reveal about the way China governs an emerging technology, and about the human costs and benefits of that approach? The questions move from surface to structure, and the analysis follows the same order.

4. Method

This is a qualitative case study built on document analysis. The corpus has four parts. The first is policy material: the artificial intelligence provisions of the 15th Five-Year Plan, national industrial-policy announcements, and the security advisories issued by the Ministry of State Security and the national computer-emergency response body. The second is corporate material: product launches,

subsidy terms, and platform announcements from Tencent, ByteDance, Baidu, and several hardware vendors. The third is community material: the discussion, tutorials, and shared jokes that circulated in open-source forums and on social platforms during the wave. The fourth is reported adopter accounts, the cases of individual users and small operators whose experience was documented in trade and general press.

These sources were read against one another to identify recurring patterns, tensions, and turning points, and then organized around a small set of analytic categories that the material itself suggested. The approach is interpretive rather than statistical. Where numbers appear, they are illustrative figures drawn from the source material, not measurements from a survey the author conducted, and they should be read as indicative of scale rather than as precise estimates. The aim is a coherent account of a phenomenon, tested against multiple kinds of evidence, not a hypothesis test.

5. Findings

5.1 Who adopted, and why

Four groups stand out in the adopter accounts, and they have less in common on the surface than their shared behavior suggests. What unites them is a fear of being left behind and a hope of a shortcut. Table 1 sets out the four profiles and the motive that recurs in each.

The white-collar case is the clearest. A worker earning a middle income hears that the firm intends to cut costs through automation, and reads the arrival of agent software as a personal deadline. The retired engineer who sits in a forum at three in the morning is chasing a different thing, a confirmation that he still counts. The parent researching a tutoring agent late at night is buying insurance against a future she cannot picture. The micro-entrepreneur is doing plain arithmetic: a cheap machine that answers customers around the clock is cheaper than the people it replaces. Different lives, one underlying posture, which later sections name and unpack.

5.2 Two technical paths: the scalpel and the bulldozer

The Chinese wave becomes legible only against a comparison. In the West, the leading agent tools follow what can be called a conservative path. Anthropic's Claude Code is the clean example: it is positioned as an assistant for professional programmers, it works through the command line, it binds tightly to one model family, and it is

deliberately fenced in, confined to the development environment and specified directories so that its reach across the rest of the machine stays small. That restraint fits enterprise compliance and keeps the attack surface narrow, and it comes with a predictable subscription price.

China's OpenClaw ecosystem took the opposite line. The tool was pitched not as an assistant but as a near-total stand-in, a single surface for running the tasks of a working life across many applications. Its users were not mainly engineers but operators looking for an edge: e-commerce sellers, small-business owners, and office workers with no technical background. It was model-agnostic, routing cheaply to low-cost domestic models through relay interfaces, and vendors wrapped it in graphical front ends so that non-technical users could run it without touching code. It also asked for high system privileges and reached into email and even online-banking pages, which made it powerful and left it badly exposed to prompt injection. Table 2 lays the two paths side by side. The contrast can be put in a sentence. The West is polishing a scalpel, precise and expensive to wield. The Chinese market rebuilt it into a bulldozer, blunt, cheap, and without much of a brake. Neither is simply better. A scalpel has a high barrier and narrow reach; a bulldozer covers ground and does not stop cleanly. In China's conditions the bulldozer was always going to draw the crowd.

5.3 Four tensions under one driver

The wave held four tensions in suspension at once, shown in Figure 1. The first is between mass adoption and official prohibition: while people queued to install the tool, state security bodies banned it outright from government offices, state-owned enterprises, large banks, and classified networks. The result is a collective bind in which everyone fears missing out and everyone knows the risk is real. The second tension is between the promise of open source and the reality of platform capture, as a free framework was steadily absorbed into the products of Tencent, ByteDance, and Baidu. The third is between grassroots invention and corporate harvesting: small operators used the tool to compete with large firms, while those firms subsidized adoption, built the ecosystem around it, and then folded the innovation back into their own platforms. The fourth is between the worship of efficiency and the alienation of labor, in which the same tool that makes a person feel like a one-person company also reduces that person to a supervisor of prompts. These are not four separate stories. They are four faces of one pragmatic drive

to use whatever works, whatever the cost, and the later discussion returns to that drive by name.

5.4 Policy on a dual track

The state did not speak with one voice, and the two voices were not in conflict so much as in division of labor. At the top, national strategy embraced the technology without reservation. The 15th Five-Year Plan refers to artificial intelligence dozens of times and sets an expansive adoption target for the end of the decade, with supporting programs to build the industrial datasets that agents need. Alongside that embrace sat a firm security line. The Ministry of State Security warned that OpenClaw's appetite for system control and its weak default settings made it an easy vehicle for attack, and that compromised office machines could be pressed into service as a botnet. Stolen credentials, the advisories noted, were already changing hands on underground forums.

The most revealing action was local rather than national. City and district governments competed to attract agent developers with a spread of incentives, an improvised federalism in which each jurisdiction tried a different formula. Table 3 collects the headline offers, and Figure 4 compares their advertised ceilings.

Two things follow from this design. The center sets boundaries the localities may not cross, and within those boundaries the localities are free to experiment and compete. The country becomes a set of parallel laboratories running slightly different versions of the same policy, and the center can watch which version works before committing to any of them.

5.5 Corporate strategy and the value flywheel

The platforms did not passively benefit from the wave; they engineered it. Tencent's sequence is representative. First it manufactured a spectacle, the free-installation events that produced the images of crowds. Then it converted that attention into compute demand, directing deployments onto its own lightweight cloud servers. Then it closed the loop, releasing a native framework and an enterprise assistant and folding an agent into WeChat, where more than a billion users could reach it through ordinary language. The through-line is plain: a cloud provider with idle server capacity needs a reason for people to consume it, and a hungry agent framework running around the clock is an excellent reason.

ByteDance and Baidu ran parallel plays, raising free interface quotas and launching one-click deployments to capture users before the open-

source version could take root. The wider structure is a self-reinforcing cycle, drawn in Figure 2. A free framework drives cloud consumption; local subsidies pull in users; users burn API tokens; the super-apps close the ecosystem; usage data flows back to improve the models; better models attract more users. Everyone in the loop has a reason to push it faster.

This is also the setting for what I call the carrot of algorithmic upgrading, a mode of governance that works by incentive rather than command. Three carrots dangle in view. The first is the instant payoff in efficiency, the seller who cuts a monthly staffing bill from the equivalent of several thousand dollars to a few dozen, which quietly binds the user to a particular API ecosystem. The second is the promise of moving up, encoded in subsidy programs that signal who is early and who is late. The third is a story about worth, the reassurance that anyone anxious is not obsolete, only mid-upgrade. The carrot does not forbid refusal. It makes refusal hard to imagine.

5.6 The economics underneath

None of this would have scaled without a change in the cost of machine reasoning. Domestic model interfaces ran roughly an order of magnitude cheaper per unit of text than the leading overseas model, and the cheapest domestic reasoning model undercut a leading Western model by a wide margin again. Figure 3 shows the comparison. When the marginal cost of a unit of automated thinking falls close to zero, the economic logic of automating repetitive office work stops being a forecast and becomes an accounting fact. A specific case circulated widely: a cross-border seller who had paid for three outsourced staff to handle inquiries and order exceptions replaced them with an around-the-clock agent and cut the monthly bill from roughly twenty thousand yuan to a few tens of dollars in interface fees, with output holding steady.

5.7 China as a governance testing ground

Put the pieces together and a proposition emerges: China became the world's testing ground for agent governance, the place where the technology met the population at scale before the rules were settled. Four features made this possible. A flexible data regime let designated zones move before national legislation caught up. A regulate-after-launch posture let rules follow the technology within weeks rather than years. A very large user base turned any national rollout into a real-scale stress test. And the central-local division of labor spread trial and error across many jurisdictions at once.

Table 4 places this model beside the two main alternatives. The appeal of this reading should not obscure its cost. When a society tries first and regulates second, the bill for the trial is paid by the people in it, in leaked data, in the risks the security bans imply, and in the older pattern of cleaning up after the fact. China's position as a testing ground is at once an advantage and a liability, and the liability lands on ordinary users.

6. Discussion

Three concepts help carry the findings past description. The first is radical pragmatism, the cultural engine of the wave. It has no loyalty to technical pedigree. If a result can be reached by stitching a Western open-source framework to the cheapest domestic compute and a national super-app, the seam is acceptable and the bugs get fixed while the thing runs. Run first and optimize later is the working rule, and it stands in contrast to a Western preference for getting the design right before release. This disposition has deep roots in a practical strain of Chinese economic culture and in two decades of internet products built on rapid iteration.

The second concept is algorithmic survivalism, which names the posture shared by the four adopter groups. This is not ordinary technology adoption. It is a defensive move made under pressure, in which a person reshapes themselves into a form that the prevailing algorithms will keep rather than discard. It shows up as emergency reskilling, when the worker whose spreadsheet skills are now automated retrains as a prompt engineer inside a few months. It shows up as the compression of time, the retiree who keeps forum hours because the technology does not sleep. And it shows up as a defensive move made on a child's behalf, tutoring bought less to make a better person than to avoid raising an obsolete one. The cruelty of the posture is that it presents an obligation in the costume of an opportunity. The user is not choosing to embrace the tool so much as proving they still have value worth extracting.

The third concept is digital narcotization, a collective coping response to anxiety that has grown too heavy to face directly. On this reading, activity becomes a way to avoid a harder truth. The parent who researches a tutoring agent at one in the morning may be less engaged in planning than in performing action, a ritual that says: as long as I am still learning, I have not been left behind. Producing data for the system becomes a kind of proof of existence, and the shared greeting that spread with the tool functions as a password among the anxious, a way to confirm that no one is afraid alone. The

danger of this response is that it helps people get used to a problem instead of solving it. Immersed in the game of lobster farming, few users paused to ask who wrote the rules of the game and who stood to win it. Table 5 draws the two underlying cultural stances together.

Read together, these three ideas describe a system that does something subtle. It does not coerce participation. It arranges conditions so that people govern themselves toward the outcome the system prefers, which is a more durable form of control than any order could be [4]. The classical theory of

governmentality anticipated exactly this move, in which power operates by shaping the field of possible action rather than by direct command [5]. What is new is the speed and the medium. The same mechanisms of deskilling and labor reorganization that industrial analysts described decades ago are now running through software that reorganizes mental work in months rather than years [6], [7]. Whether that acceleration is a liberation or a new form of dependence is not a question the data settles, and the honest answer is that it is visibly both at once.

Adopter group	Situation	Dominant motive
Mid-career white-collar workers	Salaried, facing age-based workforce pressure	Survival: avoid being cut in the next efficiency drive
Retired technical professionals	Out of the workforce, active in open-source forums	Self-worth: proof of continued relevance
Anxious parents	Researching AI tutoring for young children	Positional fear: the child must not start behind
Grassroots micro-entrepreneurs	Running one-person online businesses	Leverage: a low-cost stand-in for hired staff

Dimension	Western path (e.g. Claude Code)	Chinese path (OpenClaw ecosystem)
Product framing	Bounded assistant for specialists	Near-total stand-in for a working life
Target user	Professional software engineers	Operators, sellers, non-technical workers
Model coupling	Tightly bound to one model family	Model-agnostic, routed to cheapest option
Interface	Command line	Graphical front ends over the framework
Permissions	Fenced to the development environment	Broad system access, reaches email and banking
Business model	Fixed subscription	Free framework plus per-use token fees




Figure 1. Four tensions in the OpenClaw wave, organized around a single pragmatic driver.

Locality	Headline incentives
Longgang, Shenzhen	Up to 2m yuan for open-source code; up to 10m yuan direct investment for early

Locality	Headline incentives
	ventures; three-year free compute quota
Xuhui, Shanghai	50% R&D tax rebate for finance and gaming teams; priority access at compute centers
Wuxi High-Tech Zone	Up to 1m yuan electricity subsidy from the compute center; three-year rent exemption
Xiaoshan, Hangzhou	Annual 20m yuan fund reimbursing 30 to 50% of API token costs for domestic models
Chengdu	Up to 100k yuan settlement grants; free spare supercomputing capacity for small teams




Figure 2. The ecosystem value flywheel linking framework, compute, policy, users, and data.




Figure 3. Indicative per-unit inference costs cited during the wave (USD per 1,000 tokens).




Figure 4. Advertised ceilings of selected local incentive programs (million yuan).

Governance model	Representative instrument	Posture	Fit with agent software
European Union	EU AI Act	Precaution first, tiered regulation	Cautious; slows fast adoption
United States	Executive guidance	Industry self-regulation, light touch	Fragmented; hard to scale uniformly
China	Sandbox experiment	Launch, observe, then regulate	Close fit for rapid diffusion

Cultural dimension	Western paradigm	Chinese OpenClaw paradigm
View of technology	Elegant design	Design that works
Iteration philosophy	Perfect before release	Fix while sprinting
Standard of success	Technical leadership	Market share
Risk preference	Avoid risk	Accept uncertainty
Diffusion pattern	Radiating from an elite	Spreading from the grassroots

7. Conclusions

The wild phase of the wave will not last. A period built on command lines, third-party plugins, and cheap interfaces stitched together carries too much outage and security risk to become the steady state. As regulators press mandatory security sandboxes and the platforms roll out their own native agent services, the fever will cool and the work will move into regulated, enterprise-grade channels. That shift is not a failure of the open-source moment. It is what the open-source moment was for. In a few crowded months, OpenClaw walked tens of millions of ordinary people through their first real encounter with autonomous software, and it did so at a speed and cost that no careful program could have matched.

The episode carries a general lesson about how technologies spread. Diffusion is never only a technical matter; it is social, economic, and cultural at the same time. The Chinese wave exposed a

widespread anxiety about the future and, in the same motion, a striking capacity to adapt to it. The Western scalpel is still being sharpened. The Chinese bulldozer has already crossed the field. Which approach is wiser may have no clean answer. What is certain is that the change has happened and will not reverse, and the useful question is no longer whether to work with agents but how a person finds a place in the new arrangement, as a director of these systems rather than a thing they replace.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

- [14] State Council of the People's Republic of China. (2017). New Generation Artificial Intelligence Development Plan. Beijing: State Council.

References

- [1] Wooldridge, M. (2009). An Introduction to MultiAgent Systems (2nd ed.). Chichester: John Wiley & Sons.
- [2] Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T., & Fritz, M. (2023). Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security, 79-90.
- [3] Shifman, L. (2014). Memes in Digital Culture. Cambridge, MA: MIT Press.
- [4] Marcuse, H. (1964). One-Dimensional Man: Studies in the Ideology of Advanced Industrial Society. Boston: Beacon Press.
- [5] Foucault, M. (1991). Governmentality. In G. Burchell, C. Gordon, & P. Miller (Eds.), The Foucault Effect: Studies in Governmentality (pp. 87-104). Chicago: University of Chicago Press.
- [6] Braverman, H. (1974). Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century. New York: Monthly Review Press.
- [7] Fuchs, C. (2014). Digital Labour and Karl Marx. New York: Routledge.
- [8] Srnicek, N. (2017). Platform Capitalism. Cambridge: Polity Press.
- [9] Zuboff, S. (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. New York: PublicAffairs.
- [10] Gillespie, T. (2018). Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media. New Haven: Yale University Press.
- [11] MacKinnon, R. (2011). China's networked authoritarianism. Journal of Democracy, 22(2), 32-46.
- [12] Pasquale, F. (2015). The Black Box Society: The Secret Algorithms That Control Money and Information. Cambridge, MA: Harvard University Press.
- [13] Geertz, C. (1963). Agricultural Involution: The Processes of Ecological Change in Indonesia. Berkeley: University of California Press.