JIE GONG, JIN LI, AND FEI PU

The Double-Edged Sword of Mandatory AI Adoption at a Leading Medical Device Firm

Company Background

MedTech Corp (a pseudonym), established in 1999, is a leader in China's medical device industry. Positioned as a provider of full-lifecycle solutions for cardiovascular disease, the company operates across multiple segments, including medical devices, pharmaceuticals, and medical services/health management. It ranks among the top global players in the cardiovascular device market and is one of the few Asian enterprises in the top tier. Domestically, it holds a dominant position in niche areas such as structural heart disease treatments and vascular implants.

As of 2024, the company generated approximately RMB 6 billion in revenue and employed nearly 9,000 people. Its workforce is heavily weighted toward production (over 3,500 staff) and sales (over 2,000 personnel), supported by a strong R&D team of over 1,000 engineers. The organization comprises dozens of subsidiaries and affiliated companies, coordinating a vast network of R&D and production entities.

The Challenge: Deploying AI Assistant

To drive efficiency, the company launched an internal AI assistant. Built on the DeepSeek pre-trained models, the system was fine-tuned with a proprietary enterprise knowledge base containing technical documents, customer service records, and regulatory archives.

The tool was designed to address specific bottlenecks in high-value workflows. For international sales, entering new markets traditionally required days of fragmented research. AI could generate comprehensive market research reports in a few minutes (a 90% reduction in time) and assist in drafting localized business emails (raising response rates from 30% to 55%). It also helped domestic sales teams analyze hospital department needs and generate pitch scripts, significantly reducing sales preparation time. R&D engineers could use the tool to rapidly retrieve competitors' technical parameters during the project initiation phase, replacing "needle-in-a-haystack" manual searches with structured summaries.

Despite these demonstrated benefits, organic adoption was low. Before management intervention, usage logs showed most employees were inactive, with monthly usage counts

hovering near zero. Accustomed to traditional workflows, the workforce did not voluntarily integrate the new technology, creating a "last mile" problem in technology transformation.

The Intervention: Mandatory Usage Targets

Faced with low adoption, management decided to force the issue. On March 31, 2025, the CEO's office issued a formal notice mandating use of the AI platform for all non-frontline employees (i.e., those not working directly on the production line). The directive explicitly established a strict "monitoring and assessment" framework based on backend usage logs.

The core requirement was a minimum of 200 effective queries per month. The policy introduced a graduated enforcement scheme to ensure compliance: employees falling below 90% of the target (i.e., fewer than 180 queries) for the first time would receive a formal warning. Repeated non-compliance triggered an RMB 200 fine.

While the policy encouraged genuine feedback, employees' immediate focus became hitting the quantitative target to avoid disciplinary action.

## Results: Gaming, Distortion, and Performance

The usage pattern following the intervention provides a clear example of Goodhart's Law: when a measure becomes a target, it ceases to be a good measure.

Under the strict 200-query regime, employee behavior shifted toward "gaming" the system rather than genuine adoption. Usage logs revealed bunching: a large share of employees recorded exactly 200 queries, just enough to avoid penalties. Timing patterns also suggested limited integration into daily work. Usage remained low early in the month and surged dramatically in the final days, indicating that employees were rushing to meet a quota rather than using the tool for ongoing work needs.

Employees also complained about the time burden of generating 200 queries. In response, the company adjusted the policy in July, lowering the target to 100 queries per month. After this change, the bunching point shifted to the new 100-query target, and the intensity of the end-of-month surge flattened. While the total number of queries dropped, the quality of interaction improved: the share of repeated queries (where employees ask the same question multiple times to drive up counts) fell, and the relevance of queries to actual work tasks increased.
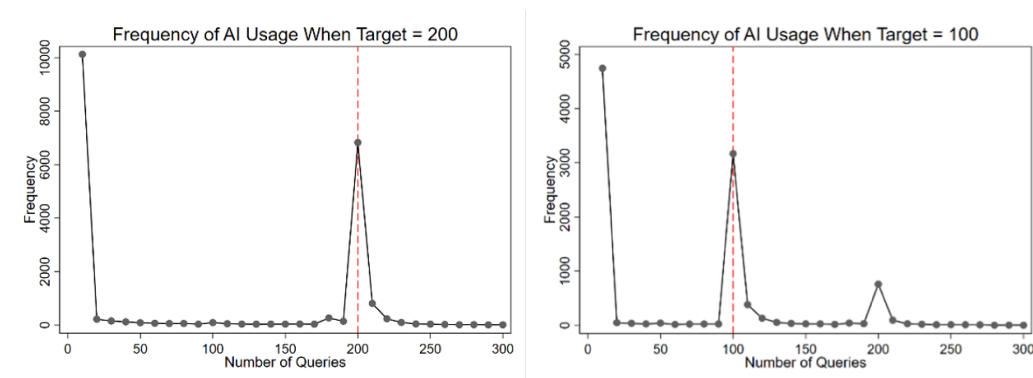
MANDATORY AI ADOPTION
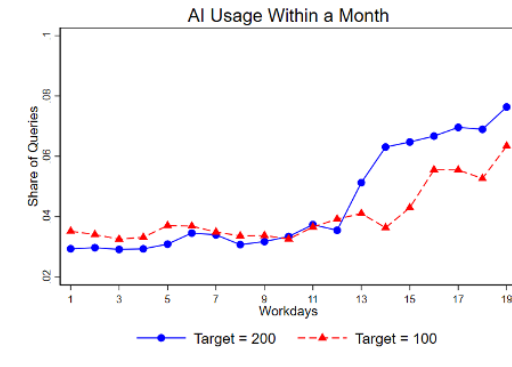
The ultimate test of the policy was its impact on employee performance. An analysis of a sales-personnel subsample revealed a stark contrast between gamed and genuine usage.

A high total number of queries was associated with lower sales performance. This negative correlation was driven by gaming behaviors—low-quality, repetitive, or rushed queries that consumed employee time without adding value. This suggests a significant opportunity cost: these queries not only fail to generate useful knowledge but also crowd out productive work time, resulting in tangible lost sales.

In contrast, when the analysis isolated work-relevant queries, the relationship flipped: higher-quality usage correlated with better performance outcomes. This indicates that the AI tool itself is a valuable asset with the potential to enhance performance when deployed correctly. Consequently, the primary challenge is not the technology itself, but how to direct employees to use it effectively instead of gaming the system.

The implementation of the 200-query mandate successfully drove immediate AI adoption but also introduced gaming behavior. The pattern on performance also raised internal debates about whether quantity-based metrics are the right tool for driving digital transformation. Management faces the challenge of designing an effective incentive structure that can encourage and sustain genuine, high-quality engagement to realize the potential of new technologies.