



SOCIETAL THRESHOLDS IN AI ADOPTION

ABOUT ZKAI

ZKAI LLC is an AI advisory and research firm founded by Zack Kass, grounded in his expertise as a leading voice in artificial intelligence. The firm spans three core areas: executive speaking, strategic advisory, and research. Zack delivers keynotes that make complex AI topics accessible to business and policy audiences, while advising executive teams on AI strategy, product direction, and organizational readiness. These experiences inform ZKAI's research, which focuses on practical, domain-relevant insights. Together, these functions create a feedback loop that keeps ZKAI's work applied, timely, and impactful.

About this Paper

This white paper is aimed at business leaders who are exploring strategic innovation and the integration of artificial intelligence (AI) within their organizations. From a product-development standpoint, a nuanced understanding of Societal Thresholds can facilitate increased consumer adoption and market penetration. Concurrently, from a managerial perspective, recognizing these Societal Thresholds can drive internal acceptance and implementation of AI solutions, thereby enhancing organizational productivity and operational efficiency.

Grounded in ZKAI's experience advising executives and analyzing past innovation waves, the paper offers a clear model for understanding how these thresholds operate. It is purposefully designed for non-technical readers, with defined terms, streamlined language, and visual frameworks that make complex ideas digestible. The aim is to help leaders anticipate adoption challenges and translate AI potential into strategic advantage.

Core Thesis: The greatest barrier to AI adoption is no longer technological; it's societal. While the scientific frontier advances rapidly, informed by the work of a few scientists, the frontier of adoption moves at a much slower pace, shaped by collective human dynamics, culture, trust, and perception.

What **will** a machine do?" is **no longer** simply answered by "what **can** a machine do?"

We present a strategic framework for understanding and navigating societal thresholds in AI deployment. We introduce a novel concept: a model for understanding the origin of societal thresholds. Organizations that master these dynamics will achieve lasting competitive advantages by building social acceptance as a strategic capability alongside technical innovation.

I. Introduction & Context

In the accelerating age of artificial intelligence (AI), technological capability alone no longer determines the trajectory of innovation. Accordingly, progress hinges on a less tangible but equally decisive factor: the readiness of society to accept and integrate these innovations into daily life, institutions, and business systems. This white paper introduces and coins the term **Societal Thresholds**.

Societal Thresholds: The constraints on the adoption of AI systems not due to technical limitations, but to social, cultural, or regulatory barriers that inhibit full utilization of what AI can already accomplish.

Crucially, these thresholds are not static or immovable. They vary across technological innovations and even more granularly across individual applications within a single innovation. For instance, while AI-generated summaries may be seamlessly adopted in communication platforms, AI-assisted decisions in clinical diagnostics or legal rulings are met with hesitation, despite displaying similar levels of technical maturity. Such variation emphasizes the need for a nuanced framework that explains the origination of these thresholds.

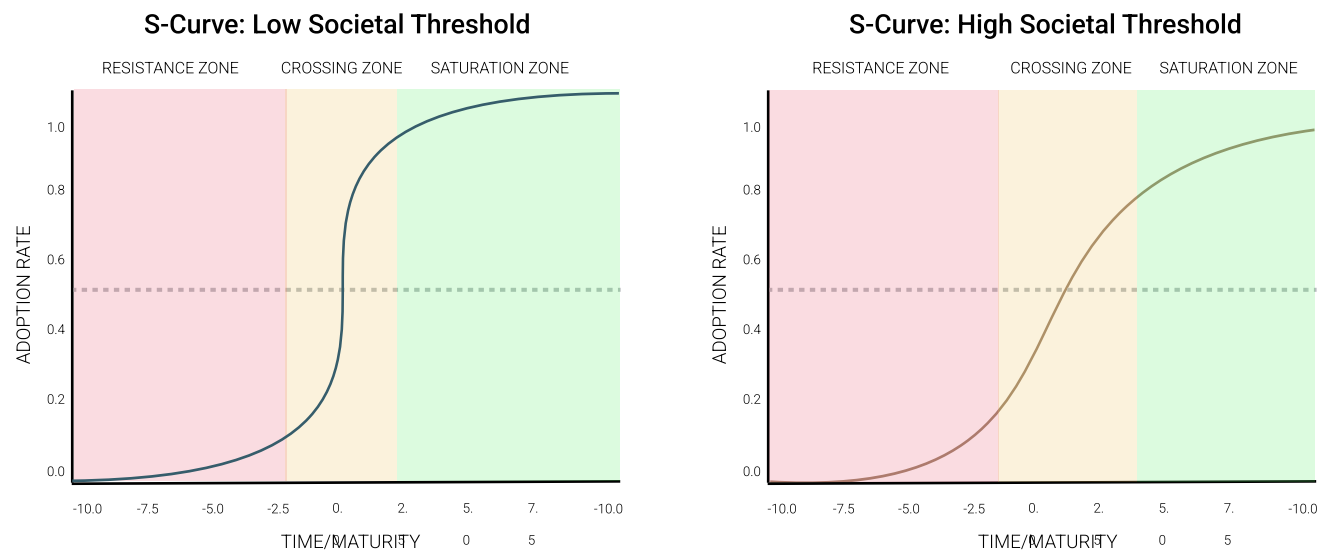
Our model for the origin of Societal Thresholds defines them as rooted in the cascading relationship where human biases shape social attitudes, which in turn become embedded in regulatory action.

Developing fluency in Societal Thresholds has become a core tactic for business leaders navigating the AI transformation. This mastery requires leaders to embody two essential traits: open-mindedness and adaptability. These qualities are not fixed attributes but cultivated abilities. Open-mindedness initiates a self-reinforcing feedback loop in which leaders who are willing to explore AI's possibilities tend to educate themselves more deeply. This, in turn, increases their capacity to adapt in how they manage their own roles and in how they restructure systems around them. This loop accelerates learning and reduces friction to produce value far beyond what technical implementation alone can achieve.

In this context, the adoption of AI is not merely a technical deployment challenge but a socio-organizational inflection point. Organizations that thrive will be those that recognize Societal Thresholds early, understand their underlying dynamics, and approach them with curiosity and empathy to inform strategy. Our model is grounded in this philosophy and aims to equip leaders to respond to resistance and transform it into momentum for innovation.

II. Impact of Societal Thresholds on S-Curves

The trajectory of any technological adoption curve (often visualized as an S-curve) is fundamentally shaped by societal thresholds. Lower thresholds of acceptance and trust lead to rapid uptake, while higher thresholds slow adoption, compressing or stretching the curve accordingly. These thresholds reflect deep-seated cultural, ethical, and psychological readiness for change, making them the invisible architecture behind every visible growth pattern. While the S-curve offers strategic value for forecasting, planning, and timing, it is only through understanding Societal Thresholds that the curve becomes truly actionable, revealing not only when adoption will occur, but the conditions that spark it. This relationship is reciprocal: analyzing the shape of the curve can surface disparities in adoption readiness and inform the calibration of ethical, regulatory, or cultural interventions needed to close the gaps.



Case Study: AI-Enhanced Irrigation Tools

Across developing regions, agriculture faces chronic challenges: scarce water, poor yields, erratic climate, and limited agronomic expertise (Shepherd et al., 2020). Yet transformative solutions do not always require massive infrastructure or policy overhauls—they can scale rapidly when adoption barriers are low (World Bank, 2021). AI-based mobile irrigation advisors, developed by organizations like Digital Green and Precision Agriculture for Development (PAD), exemplify this principle. These tools operate through SMS and mobile apps, using AI to analyze weather, recommend irrigation, suggest crop rotations, and deliver alerts in local dialects with minimal infrastructure or training required (Gill et al., 2022).

The result is a distinct S-curve: unlike many technologies that stagnate in early resistance phases, these tools accelerate into widespread use due to their low societal threshold—defined by cost, usability, offline functionality, and linguistic accessibility (PAD, 2023). Adoption occurs at the individual level, yielding 10–40% increases in productivity and more efficient water use (Agrosmart, 2022). This case demonstrates that successful diffusion depends as much on contextual alignment as technical sophistication, illustrating the unique impact of Societal Thresholds on the S-curve.

Case Study: AI in Healthcare

AI systems have shown exceptional promise in medical diagnostics, as they match or even outperform clinicians in detecting diseases such as cancer, diabetic retinopathy, and skin lesions (Esteva et al., 2017; McKinney et al., 2020). Yet despite proven technical capabilities, the healthcare sector has struggled to move beyond the early stages of adoption because of a high Societal Threshold defined by regulation, liability, and institutional trust.

Unlike consumer-facing technologies, AI in healthcare must navigate strict oversight from agencies like the FDA, contend with legal ambiguity around malpractice, while earning the confidence of clinicians. Many AI tools function as “screeners,” augmenting doctors rather than replacing them due to the ethical risks of opaque decision-making. Moreover, integration into existing hospital systems requires significant technical and cultural adaptation, which many institutions are not yet equipped to manage (Topol, 2019).

This has shaped an elongated S-curve, where progress is bottlenecked in the Resistance phase. Despite strong performance in research trials, adoption remains fragmented, with most use cases confined to administrative tasks or controlled environments. Until frameworks for explainability, responsibility, and clinician integration are in place, AI in healthcare will remain technically advanced but socially encumbered.

III. Origins of Societal Thresholds

Societal Thresholds represent tipping points where collective comfort, trust, and acceptance either accelerate adoption or create resistance barriers. They are a reflection of three universal truths rooted in deeply human factors:

Humans Desire Control

The fundamental human need for control arises from deeply rooted psychological, ethical, and practical considerations. Throughout history, control has been associated with predictability and certainty, which the human mind interprets as security. Consequently, the preoccupation with maintaining control over new technologies reflects an evolutionary goal to seek safety and minimize uncertainty in the environment. Thus, when new technologies are introduced, concerns arise from a perceived displacement of human security and safety.

Previously, this desire for control permeated throughout internet and software adoption cycles. As early as the 1980s and 90s, when large software applications were introduced, the public began to worry about bugs leading to catastrophic failures related to airplane control, power grids, and weapons systems. Through increased exposure and education, public fears were assuaged as software applications now dominate daily personal and professional workflows.

AI development is following a similar trajectory. However, unlike software and web adoption, AI represents a deeper unknown. For instance, whereas software produces prescribed results, AI outputs can vary greatly across differing prompt techniques, LLM models, and fine-tuning methods. Even so, the same model can produce different results solely due to its embedded probabilistic roots. This directly undermines the human want for control. Certainty is not a feature of AI outputs.

High Threshold Scenarios	Low Threshold Scenarios
<div><p>Autonomous Vehicles</p><p><u>Control Loss:</u> Autonomy in potentially fatal automobile incidents</p><p>68% of Americans are afraid of self-driving cars, up from 55% in 2022. This is a 13% increase and the largest annual rise since 2020 (AAA).</p></div>	<div><p>Predictive Maintenance Systems (PdM)</p><p><u>Control Retained:</u> Final decision authority remains with humans</p><p>PdMs use AI sensors to forecast equipment failures while humans focus on quality improvement as opposed to routine maintenance and upkeep tasks.</p></div>
<div><p>AI Hiring Systems</p><p><u>Control Loss:</u> Decision-making regarding career outcomes and personal impact.</p><p>Employees express a lack of trust in resume screening and attrition algorithms (Penfold, Chun, Cremer; 2025).</p></div>	<div><p>Customer Support AI Agents and Chatbots</p><p><u>Control Retained:</u> Human agents are able to handle complex cases while repetitive tasks are offloaded</p><p>85% of customers prefer AI chatbots over human agents when interacting with communications service providers, citing faster outcomes and convenience (CSG, 2025).</p></div>

Essentially, the control displacement of a decision is a constant; thus, when high impact decisions are outsourced to AI, human control and identity is displaced. This becomes increasingly apparent examining discrepancies between high and low threshold scenarios. In many low threshold scenarios, not only is human control retained, but they are often able to reallocate their cognitive resources towards more complex issues. For employees, this provides a higher level of control, raising their overall job satisfaction level (Monteiro, 2023). In high threshold scenarios, the stakes are higher, as the power given to the AI system deals with personal life outcomes such as individual health or career paths.

III. Origins of Societal Thresholds


Technology Advances Faster than Society Adopts

Technology is advancing at an exceptional speed. Previously, Moore’s Law dictated the rate of progression as noted by the Intel Co-Founder Gordon Moore. He claimed, that the number of transistors on a chip would double every 2 years, leading to exponential increases in computing power and capabilities. This development rate persisted for years, leading to lower compute costs and supporting the transition from desktop PCs to smartphones for cloud computing. Since then, Moore’s Law has adopted a new interpretation. While physical limits regarding the number of transistors on chips have reached an asymptote, progress in parallel compute and software development has expedited technological progress in recent years, despite the formal death Moore’s Law.

Therefore, technology is developing quickly not only on the hardware frontier but with ricochetting effects into software and AI development. Compounded together, the scientific frontier as a whole is racing ahead of societal and business frameworks. Shifts in technology move at a rate independent of cultural attitudes and regulatory policies. Even when cultural attitudes and regulation adapt to a new innovation, business adoption often lags further behind.

Accounting for the evolving nature of Moore’s Law, the gap between technological advancements and societal/ business adoption is increasing. For instance, the AI scientific frontier is moving at an exponential rate, with AI models achieving human-level performance on complex tasks in record time. A 2025 report by Innovation Endeavors highlights that new AI models often hold the top usage spot for only about three weeks before being overtaken by open-source rivals, indicating rapid development cycles. In contrast, business adoption is lagging; a Gallup poll reveals that only 9.2% of U.S. businesses reported using AI in Q2 2025, up from 5.7% in late 2024, showing a gradual but still limited integration into business operations.

AI SCIENTIFIC FRONTIER vs. ADOPTION FRONTIER



III. Origins of Societal Thresholds

Higher Tolerance for Human Failure than Machine Failure

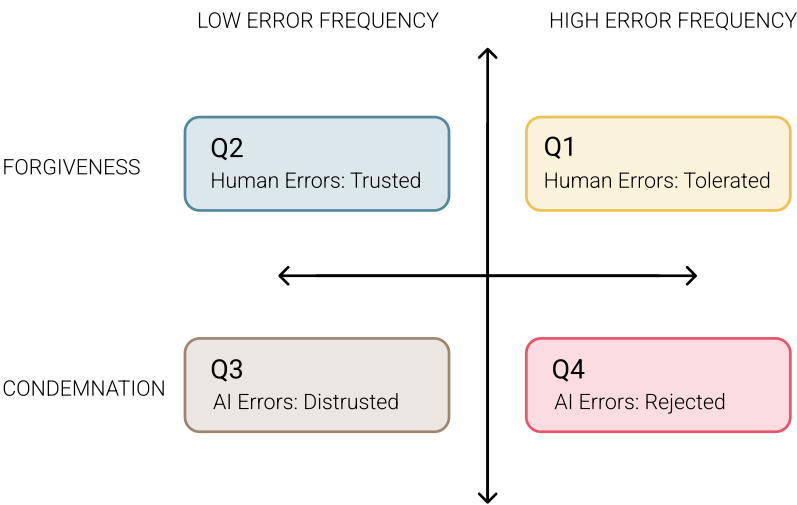
It would be reasonable to expect that societal adoption of AI might significantly trail scientific advancement, particularly if systems were demonstrably inferior to their human counterparts. However, this is not the reality. In a wide array of fields, from medical diagnostics to legal reasoning and language processing, AI systems have consistently demonstrated superior accuracy and capability when compared to human performance.

In radiology, a study published in Nature (2019) showed that a deep learning algorithm outperformed six radiologists in detecting breast cancer, correctly identifying more cases while simultaneously reducing false positives. Similarly, in legal applications, research from Stanford University found that AI tools could predict Supreme Court decisions with approximately 70% accuracy, outperforming legal experts who achieved closer to 66% (Katz et al., 2017). In natural language processing, systems like GPT-4 have demonstrated top percentile performance on standardized tests, surpassing average human scores on many metrics (OpenAI, 2023).

Despite boundless empirical evidence, a persistent perception remains: AI is viewed as less trustworthy than humans. This discrepancy reflects a cognitive asymmetry in how errors are judged. When humans make mistakes, such errors are often attributed to the natural imperfections of the human condition — an inherent and largely accepted limitation. In contrast, when AI systems err, the reaction is often condemnation, as the failure is seen as indicative of a systemic flaw.

Ironically, the technical nature of AI errors frequently allows for more straightforward remediation. While humans are capable of learning from their mistakes, the process is variable and sometimes unpredictable. In contrast, upon identification, errors in AI systems can often be corrected through precise engineering interventions, allowing for consistent and lasting improvements in performance.

This produces a paradox, wherein the more correctable system is viewed as the less reliable one, presenting a fundamental threshold in the social integration of advanced technologies. Addressing this perception gap is critical to aligning public trust with empirical capability.



The matrix expresses a key difference in societal perception between human and AI error. Assuming the same error was committed by the two parties, a human who errs often, is tolerated more than that of a machine at a lower rate. Thus, the defining factor in societal perception is the party that produces the mistake, rather than the effect of the mistake itself, proving failure asymmetry between humans and machines. This phenomenon indicates that AI products need to demonstrate trustworthiness in addition to technical capability to flourish with consumer groups.

IV. Conclusion & Application


Societal Thresholds are shaped by these three factors, encapsulating adoption dynamics. In reality, these thresholds represent barriers present in an innovation's transition from niche application to mainstream usage. Crucially, these thresholds vary significantly depending on the particular innovation or use case. Thus, adoption curves are a summation of individual Societal Thresholds across industries.

When corporations master the intricacies of Societal Thresholds, they gain strategic advantages in product development and market positioning. Therefore, by utilizing the Societal Threshold model to assess target user sentiment and concerns, companies can tailor their strategies to address threshold-produced resistance points. This approach generates a nuanced understanding of how the market operates in tandem with societal readiness levels.

Importantly, Societal Thresholds should not be viewed purely as an obstruction to innovation. They often serve as essential ethical firewalls that protect individual rights and privacy, behaving as a natural filtering mechanism. Most importantly, these firewalls ensure that the implementation of technologies undergo scrutiny before they are widely implemented. Otherwise, if AI systems are deployed prematurely in resistant industries (i.e., education, healthcare), there may be distinct negative consequences.

Additionally, circumventing these thresholds can prove counterproductive. Companies that attempt to leapfrog societal readiness often trigger backlash from later adoption groups, creating heightened resistance that ultimately slows overall adoption rates.

The framework of Societal Thresholds ultimately offers a roadmap for timely innovation by balancing the urgency of technological progress with the reality of measured adoption. Organizations that align themselves with these thresholds and address their underlying concerns will achieve greater commercial success while contributing more sustainably to technological integration. Crossing Societal Thresholds should not be treated as a race, rather as a means of building the trust necessary for lasting transformation.



Contributors

█ Zack Kass, Global AI Advisor & former Head of GTM at OpenAI

█ Montana Showalter

█ Christina Kodsí

Published: July 2025

References

- Agrosmart. (2022). How digital irrigation helps farmers save water and increase productivity. <https://agrosmart.com/en/digital-irrigation>
- Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. *BMC Medical Informatics and Decision Making*, 20, Article 310. <https://doi.org/10.1186/s12911-020-01332-6>
- American Automobile Association. (2024). 2024 autonomous vehicle survey: Public skepticism rises. AAA Foundation for Traffic Safety. <https://newsroom.aaa.com/autonomous-vehicle-survey-2024>
- CSG. (2025). Customer experience report: AI in telecom service delivery. CSG International. <https://www.csgi.com/resources/customer-experience-report-2025>
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115?118. <https://doi.org/10.1038/nature21056>
- Gill, T., Lodhi, S. A., & Oluoch, T. (2022). AI for smallholder farmers: Case studies from sub-Saharan Africa and South Asia. CGIAR Platform for Big Data in Agriculture. <https://bigdata.cgiar.org>
- Leveson, N. G., & Turner, C. S. (1993). An investigation of the Therac-25 accidents. *IEEE Computer*, 26(7), 18?41. <https://doi.org/10.1109/MC.1993.274940>
- Markoff, J. (2023, March 24). Gordon E. Moore, Intel co-founder behind Moore's Law, dies at 94. *The New York Times*. <https://www.nytimes.com/2023/03/24/technology/gordon-moore-dead.html>
- McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., ... & Suleyman, M. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, 577(7788), 89?94. <https://doi.org/10.1038/s41586-019-1799-6>
- National Geographic Society. (n.d.). Y2K bug. National Geographic Education. Retrieved July 1, 2025, from <https://education.nationalgeographic.org/resource/Y2K-bug/>
- PAD (Precision Agriculture for Development). (2023). Annual Impact Report. <https://precisionag.org/impact>
- Penfold, B., Chun, M., & Cremer, L. (2025). Algorithmic bias in HR: Trust and transparency in employee-facing AI systems. Cambridge University Press.
- Shepherd, A., Stoian, D., & Cotula, L. (2020). Harnessing digital technology for agrifood transformation. Food and Agriculture Organization (FAO). <https://www.fao.org/documents/card/en/c/cb2773en>
- Topol, E. (2019). Deep medicine: How artificial intelligence can make healthcare human again. Basic Books.