The Centre for Civic Engagement

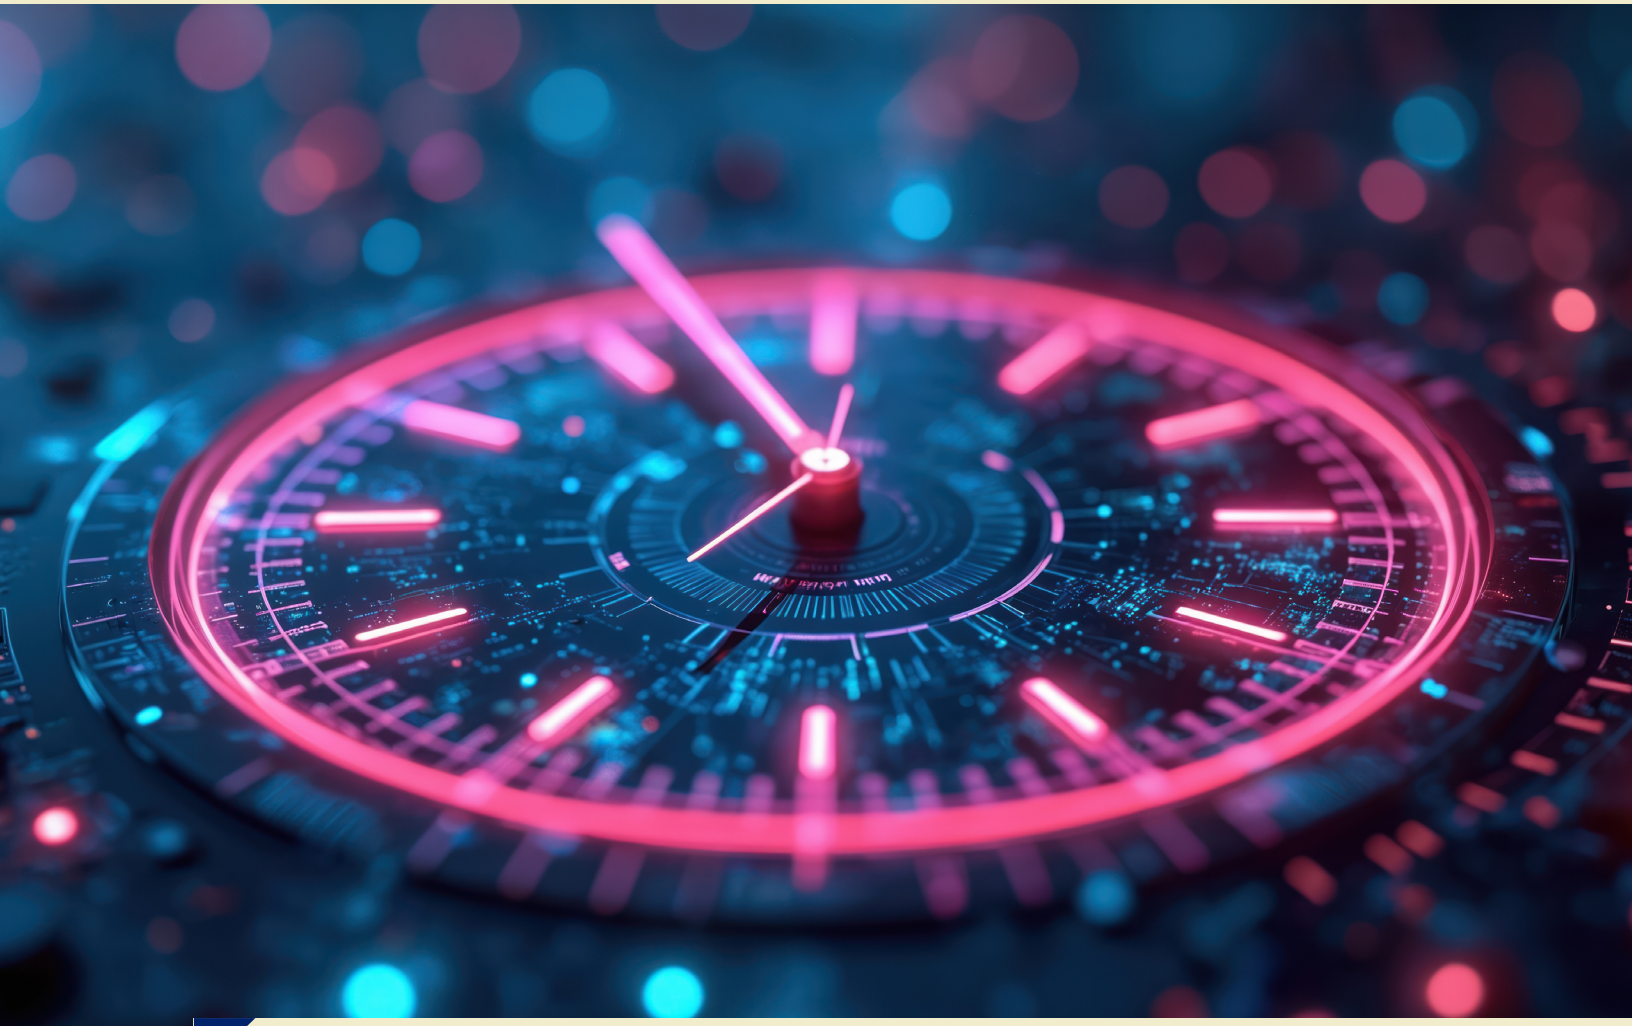

# Can We Survive the Journey to AI Abundance?

Sean Speer

This study was commissioned by the Centre for Civic Engagement. The CCE is a non-partisan Canadian charity dedicated to conducting original research on public policy issues related to Canadian prosperity, productivity, and national flourishing. The CCE's research informs an active program of policy seminars, events, conferences, and lectures all aimed at providing the policy making community with actionable insights that encourage informed decision making on issues that matter to Canadians.

cce | The Centre for Civic Engagement

# Introduction

In recent weeks, five long-form commentaries on artificial intelligence and the economy have galvanized financial markets and policy circles.

It began with a viral essay[1] by AI founder and investor Matt Shumer, who argued that we are underestimating both the pace and magnitude of AI's advance. Writing from inside the industry, he warned that recent model releases mark a step-change in capability and that large swaths of white-collar work could be displaced far sooner than most people expect.

That essay was quickly followed by four macroeconomic responses. Citrini Research, a U.S.-based market research firm, published[2] a bearish scenario entitled "Global Intelligence Crisis." Venture investor Michael Bloch countered[3] with "Global Intelligence Boom." The Kobeissi Letter, a newsletter on global capital markets, argued[4] markets are overpricing doom. And economist Tyler Cowen dismissed[5] the idea that an AI-driven world could suffer from a sustained aggregate demand shortfall.

Each scenario is internally coherent. Each rests on defensible economic assumptions. And each is trying to answer the same question: does widespread AI adoption produce abundance or instability?

But in reading them side by side, it struck me that they are largely arguing about end states. They debate whether the eventual equilibrium is prosperous or fragile. What they understate is the path between here and there.

My own instinct is that AI will ultimately be additive to human flourishing. But the transition to that equilibrium may be more economically and politically destabilizing than either the optimists or the equilibrium theorists fully appreciate. The central variable in this debate isn't just productivity, distribution, or even finance. It is time.

1 Matt Shumer (@mattshumer_), "Something Big Is Happening," X, February 10, 2026, https://x.com/mattshumer_/status/2021256989876109403?s=20.

2 Alap Shah, "The 2028 Global Intelligence Crisis," *Citrini Research*, February 22, 2026, https://www.citriniresearch.com/p/2028gic.

3 Michael Bloch (@michaelxbloch), "THE 2028 GLOBAL INTELLIGENCE BOOM," X, February 22, 2026, h ttps://x.com/michaelxbloch/status/2025712344123236418.

4 The Kobeissi Letter (@KobeissiLetter), "It's Too Obvious. What If AI Doesn't Actually End The World?," X, February 23, 2026, https://x.com/KobeissiLetter/status/2026125293149163694.

5 Tyler Cowen, "Is there an aggregate demand problem in an AGI world?," *Marginal Revolution*, February 24, 2026, https://marginalrevolution.com/marginalrevolution/2026/02/is-there-an-aggregate-demand-problem-in-an-agi-world.html.

# Is Time the Real Risk?

One way to understand why these essays disagree so sharply is that they're speaking past one another on the time horizon.

In the long run, most economists are comfortable with the idea that economies tend toward new equilibria. Technological shocks destroy some jobs and create others.[6] Prices adjust, capital reallocates and new industries emerge. Even the more pessimistic scenarios generally concede that a stable endpoint is possible.

The dispute is over what happens between now and that endpoint, and whether the transition costs are small enough to be absorbed or so large that they become economically and politically destabilizing.

This matters because AI is not just another incremental improvement. It is a general-purpose capability[7] with the potential to diffuse across white-collar workflows unusually fast. Electricity and the internet transformed economies, but adoption unfolded over decades. If AI's adoption curve is steeper—if it compresses years of organizational and technological change into months—the usual adjustment story may still apply in theory, but it may arrive too late to prevent severe turbulence in practice.

Shumer's essay is, at its core, an argument about velocity. He points to evidence that AI systems are now contributing to their own development, debugging training processes, and writing significant portions of the code used to build future models. In his telling, the timeline most people are implicitly assuming is already obsolete.

Citrini Research translates that acceleration thesis into a financial stress scenario. If capability is advancing that quickly, then labour displacement could cascade through consumer spending, private credit, and housing markets before a new equilibrium is reached. Bloch's scenario, by contrast, can be read as a theory of the equilibrium arriving quickly enough to make the transition manageable.

Cowen's intervention is different. His focus is on the coherence of the equilibrium itself. In a world producing a flood of valuable goods and services, he argues, it is hard to sustain a permanent demand shortfall.

This may be right as a matter of general economic logic, but it does not resolve the more practical question of whether societies can tolerate the route to that outcome. The central policy question, in other words, may not be "Will we get to abundance?" but rather "What happens before we do?"

[6] Daron Acemoglu and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," *NBER Working Paper Series* no. 23285 (2017).

[7] Timothy F. Bresnahan and Manuel Trajtenberg, "General Purpose Technologies "Engines of Growth?"," *NBER Working Paper Series* no. 4148 (1992).

# Can the Economy Absorb a Shock at the Top?

This time-horizon problem becomes especially acute because AI displacement is likely to land first and hardest on high-skilled knowledge workers: the people who design systems, write memos, analyze markets, build models, manage projects, and draft contracts. Basically, the administrative and professional layers of the modern economy.

Shumer's essay is explicit about this. He argues that the capability leap in coding and cognitive work is already rendering large portions of technical and professional labour redundant. Citrini Research takes that premise and pushes it forward into a macro-financial stress scenario: if white-collar productivity can be automated at scale, income compression in these cohorts is both predictable and imminent.

That has macroeconomic implications that differ from the typical recession story. In an ordinary downturn, job losses are often concentrated among lower-income or cyclically sensitive workers. The consumption impact is immediate but, in aggregate terms, dispersed across a wide base. In an AI-driven displacement scenario, the initial shock may be disproportionately concentrated in the upper-middle and upper portions of the income distribution.

This matters because higher-income households account[8] for a large share of discretionary spending—including housing upgrades, restaurants, travel, durable goods, tuition, renovations, vehicles, and professional services. Even if they represent a smaller share of total employment, they're a large share of total demand.

Citrini's concern about housing markets, private credit exposure, and financial fragility rests heavily on this composition effect: if the spending engine is disproportionately powered by high earners, income shocks at the top propagate differently than broad-based unemployment at the bottom.

There's also a second-order dynamic here that is easy to miss. High-income workers also tend to carry large fixed obligations—mortgages, childcare, tuition commitments—contracted on the basis of stable earnings trajectories. When income security is impaired, households may not collapse into destitution overnight, but they will behave defensively: savings rise, consumption falls, major purchases are delayed. The behavioural shift may lag the job loss itself, but once it eventually arrives it can be economically meaningful.

Bloch's counter-scenario effectively assumes that this transition is either short-lived or offset by rapid entrepreneurial redeployment and falling prices. The transmission mechanism is AI-driven deflation, which, in his telling, raises real purchasing power quickly enough to cushion the blow. Cowen goes even further, arguing that in a world of abundant new goods and services, aggregate demand cannot remain depressed for long because prices and incentives will adjust.

But this is precisely where the time variable re-enters. Even if an abundance equilibrium exists, the path to it may involve precisely the kind of temporary but sharp demand softening that triggers broader economic weakness. Economists can say "prices will adjust." Policymakers have to ask how long that adjustment takes—and whether the political system can tolerate the interim.

[8] "Consumer Expenditure Surveys," *U.S. Bureau of Labor Statistics,* https://www.bls.gov/cex/.

# Is the Welfare State Built for an AI Shock?

If the dislocation hits mid-career knowledge workers—people with high human capital and high earnings—the fiscal effects and policy design challenges become more complicated.

Shumer's essay is explicit that the first wave of displacement is unlikely to be factory workers or retail clerks. It's software engineers, analysts, lawyers, consultants, managers and other members of the professional cohort that modern economies have trained and credentialed for decades. Citrini Research takes that premise and models what happens when that cohort's income is impaired quickly and at scale.

But neither the optimistic nor the equilibrium-oriented responses fully grapple with what that means for the structure of the modern welfare state.

Modern welfare states weren't designed to cushion large income shocks at the upper end of the distribution. Employment insurance is capped and replaces only a portion[9] of income. It may prevent poverty, but it does not maintain a professional household's previous standard of living. Automatic stabilizers were built to smooth cyclical downturns among median earners—not to absorb structural income compression among six-figure professionals with large fixed obligations.

Bloch's scenario implicitly assumes that this mismatch is temporary. Entrepreneurial redeployment and falling prices restore real purchasing power before fiscal stress compounds. The Kobeissi Letter similarly argues that markets are underestimating how quickly deflation and new demand channels will offset dislocation. In both accounts, the adjustment happens fast enough that the fiscal architecture bends but doesn't break.

Citrini's scenario, by contrast, assumes the opposite sequencing. Income compression arrives first. Households respond by cutting consumption and other discretionary spending. Only later does a new equilibrium emerge. That ordering matters enormously for fiscal capacity.

If millions of high-income earners experience displacement or persistent income loss, the economic effects are not confined to their personal finances. These households contribute disproportionately to tax revenues, including income taxes, payroll taxes, and consumption taxes. If their earnings fall, receipts will decline. Meanwhile, spending pressures rise as more people draw on benefits and as political demands grow for more expansive transitional support.

This is the fiscal version of the same time-horizon problem: the state may be asked to do more precisely when its revenue base is impaired.

Cowen's argument is that in a world producing a flood of valuable new goods and services, aggregate demand cannot remain depressed for long. Prices will adjust and incentives realign. Income is generated somewhere in the system. That may well be correct in general equilibrium terms. But from a fiscal perspective, the question isn't whether income exists in the system. It is where it accrues and how quickly governments can tax it.

[9] Canada, Government of Canada, *Employment Insurance benefits*, Benefits, Ottawa, 2026, https://www.canada.ca/en/services/benefits/ei.html.

Which leads to an unavoidable question that the viral essays largely neglect: what becomes the tax base in an AI-intensive economy?

If labour income becomes more volatile governments may need to rely more heavily on other revenue-generating tools like consumption taxes, capital income, resource rents, or levies tied[10] to AI infrastructure and compute. Shumer hints at an economy increasingly powered by a small number of AI labs and compute clusters. Citrini worries about the financial fragility tied to that concentration. Bloch assumes the gains diffuse widely through price compression. But none fully specify how public finance adapts if the composition of income shifts materially toward capital and infrastructure.

Even a government that issues its own currency cannot ignore the political economy of taxation and legitimacy. Fiscal capacity ultimately rests on broad participation in the tax base and on public consent about who bears the burden. If displacement is rapid and concentrated among previously secure cohorts, the political response may precede the economic re-equilibration.

That's why the sequencing matters so much. Even if the arithmetic of productivity growth may be favourable, the institutional capacity to manage the transition is far less certain.

[10] Tyler Cowen, "Public Finance in the Age of AI: A Primer," *Marginal Revolution*, February 25, 2026, https://marginalrevolution.com/marginalrevolution/2026/02/public-finance-in-the-age-of-ai-a-primer.html.

# Will People Feel Better Off?

One of the most important and least examined fault lines between the boom and crisis scenarios is psychological.

Bloch's abundance story leans heavily on the distinction between nominal income and real purchasing power. If AI drives sharp deflation in services, households can be materially better off even if wages stagnate or decline. The Kobeissi Letter similarly argues that margin compression in one sector becomes cost savings for consumers elsewhere. Even Cowen's equilibrium logic rests on the idea that if goods and services flood the economy, income must ultimately be generated somewhere and prices will adjust accordingly.

As an economic proposition, this is coherent. But it raises a deeper human and political question: can you persuade millions of people that they are better off when their paycheques are falling, their career trajectories are disrupted, and their sense of professional identity is shaken?

Shumer's warning is about more than income loss. It envisions a world of redundancy in which highly trained professionals discover that the core of their expertise can be replicated in plain English by a machine. Citrini's crisis scenario imagines how quickly such a realization could translate into defensive household behaviour, falling asset values, and financial strain. Both implicitly recognize that labour income is not just a consumption stream. It's a signal of status, agency, and security.

People experience progress through wages, promotions, and the ability to convert skill into income. A world in which living standards rise primarily through lower prices rather than higher incomes would require a shift in how individuals understand advancement, security, and dignity. It would mean telling a displaced mid-career professional that cheaper legal advice, lower travel costs, and more efficient services compensate for the loss of a hard-earned role in the productive hierarchy.

That shift may eventually occur. But it's a major adjustment to the narrative that we've told people for decades. Bloch assumes adaptation and entrepreneurial redeployment happen quickly enough to cushion the shock. Cowen assumes that Say's Law logic prevents sustained demand shortfalls—in an economy producing a flood of valuable goods and services, income is necessarily generated somewhere in the system and prices adjust until that output is absorbed. Yet Citrini and Shumer raise doubts that social psychology and institutional adjustment can keep pace with these exponential capability gains over the short and medium terms.

This is where a technological transition can become socially ugly even if the long-run economic endpoint is favourable. A society can become more materially abundant while simultaneously more resentful and unstable if the transition undermines widely held expectations about how effort, skill, and reward are supposed to relate.

Abundance doesn't automatically produce legitimacy. And legitimacy, in a democratic society, is an economic variable too.

# Conclusion

It's entirely possible to believe that AI will ultimately be enormously beneficial—perhaps transformative in fields like medical research and life-extending discovery—while also believing that the near-term adjustment could overwhelm institutional capacity and generate severe social conflict.

The main risk could be a time-horizon problem. It suggests that the key question is not "boom or crisis?" but "what is the transition path, and how do we widen the bandwidth of adaptation?"

The four essays are useful precisely because they stress different failure modes and different stabilizing mechanisms. Citrini warns of feedback loops when adjustment is slow and finance amplifies shocks. Bloch reminds us that entrepreneurship and price declines can be powerful stabilizers if reallocation is fast enough. The Kobeissi Letter cautions that markets can mistake repricing for ruin. Cowen offers a necessary check against incoherent macro claims that ignore equilibrium logic.

The policy task is to treat these scenarios as stress tests. None of them are destiny. We have agency in influencing the conditions that would make each more (or less) likely.

If the main risk is transition speed, then the priority becomes strengthening labour mobility, capital flexibility, and financial resilience. If the main promise is services deflation and entrepreneurial dynamism, then the focus should be on ensuring competition, diffusion, and broad access to capital and tools.

Policy cannot determine the direction of technology. But it can shape whether a high-speed transition produces stability or fracture.

AI may represent the most powerful productivity shock[11] in modern history. Whether it produces broad-based prosperity or destabilizing volatility will depend less on the technology itself than on how quickly our institutions, labour markets, and fiscal systems adapt to it.

The real debate, then, is not about abundance versus collapse. It's about whether we can survive the journey to abundance without tearing ourselves apart along the way.

[11] "Generative AI could raise global GDP by 7%," *Goldman Sachs*, April 5, 2023, https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent.

# Bibliography


Acemoglu, Daron and Pascual Restrepo. "Robots and Jobs: Evidence from US Labor Markets." *NBER Working Paper Series* no. 23285 (2017). https://www.nber.org/papers/w23285.

Bloch, Michael (@michaelxbloch). "THE 2028 GLOBAL INTELLIGENCE BOOM." X, February 22, 2026. https://x.com/michaelxbloch/status/2025712344123236418.

Bresnahan, Timothy F. and Manuel Trajtenberg. "General Purpose Technologies "Engines of Growth?" *NBER Working Paper Series* no. 4148 (1992). https://www.nber.org/papers/w4148.

Canada. Government of Canada. *Employment Insurance benefits*. Benefits. Ottawa. 2026. https://www.canada.ca/en/services/benefits/ei.html.

"Consumer Expenditure Surveys." *U.S. Bureau of Labor Statistics*. https://www.bls.gov/cex/.

Cowen, Tyler. "Is there an aggregate demand problem in an AGI world?" *Marginal Revolution*. February 24, 2026. https://marginalrevolution.com/marginalrevolution/2026/02/is-there-an-aggregate-demand-problem-in-an-agi-world.html.

Cowen, Tyler. "Public Finance in the Age of AI: A Primer." *Marginal Revolution*. February 25, 2026. https://marginalrevolution.com/marginalrevolution/2026/02/public-finance-in-the-age-of-ai-a-primer.html.

"Generative AI could raise global GDP by 7%." *Goldman Sachs*. April 5, 2023. https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent.

Shah, Alap. "The 2028 Global Intelligence Crisis." *Citrini Research*. February 22, 2026. https://www.citriniresearch.com/p/2028gic.

Shumer, Matt (@mattshumer_). "Something Big Is Happening." X, February 10, 2026. https://x.com/mattshumer_/status/2021256989876109403?s=20.

The Kobeissi Letter (@KobeissiLetter). "It's Too Obvious. What If AI Doesn't Actually End The World?" X, February 23, 2026. https://x.com/KobeissiLetter/status/2026125293149163694.

**Sean Speer** is the editor-at-large at The Hub. He previously served as a senior adviser to former Canadian Prime Minister Stephen Harper.

cce | The Centre for Civic Engagement

centrecivicengagement.ca