

Contributors

René Lalonde

Director, Modelling and Forecasting
Scotiabank Economics
416.862.3174
rene.lalonde@scotiabank.com

Olivier Gervais


Director, Modelling and Forecasting
Scotiabank Economics
416.288.7836
olivier.gervais@scotiabank.com

Farah Omran

Senior Economist
Scotiabank Economics
416.866.4315
farah.omran@scotiabank.com

Chart 1

U.S. AI-Adjacent Investment Has Accelerated Sharply



Sources: Scotiabank Economics, BEA.

The AI Shock: How It Could Shape the Economy

- **AI is best understood as a productivity shock that works through task-level cost savings, with spillovers to labour markets, investment, asset valuations, and inflation. The scale and timing of these effects remain highly uncertain.**
- **The labour-market channel is the key uncertainty: outcomes depend on whether AI mainly substitutes for workers or complements them, and on how quickly retraining and new task creation offset displacement.**
- **Scenario analysis is therefore the right framework. Our upside cases consider stronger productivity growth with either limited labour-market disruption or a more uneven adjustment path, while the downside centres on a repricing if AI gains fail to materialize on the expected scale or timeline.**
- **In the upside scenarios, stronger productivity lifts the level of GDP in both the U.S. and Canada, though the inflation path depends on the transition as productivity creates excess supply. An initial AI investment build-out can speed up the adjustment of demand and delay disinflation before supply gains begin to dominate.**
- **The main downside risk is that expected productivity and earnings gains do not arrive on time, triggering an equity repricing, tighter financial conditions, and weaker demand.**
- **Across scenarios, one conclusion is consistent: AI is likely to be disinflationary over time, whether through stronger supply in the upside cases or weaker demand in the downside case.**

Artificial intelligence has rapidly become a central macroeconomic question.¹ Its potential to reshape productivity, labour markets, and asset valuations places it alongside past general-purpose technologies such as electrification or the internet. Yet, unlike many earlier technological waves, the speed of diffusion, the breadth of potential applications, and the scale of market expectations introduce an unusually high degree of uncertainty.

While we are still in the early stages of the AI cycle, AI investment is already well underway, particularly in the U.S., which has led the build out of AI-enabling infrastructure. U.S. spending on computer equipment and data-centre structures² has accelerated sharply in recent years (chart 1). Adoption is harder to measure, as evidence remains more limited and uneven, relying largely on survey data that vary by AI definitions and scope of use. Nevertheless, recent survey evidence suggests material **diffusion** of AI in the U.S., whereas in Canada, operational use **remains** at an earlier, more experimental stage.

At this stage, both the magnitude and timing of AI's future economic impact remain unclear with estimates of productivity gains varying widely across studies. Measured effects on productivity are still difficult to identify in real time, and the transmission to labour markets, inflation dynamics, and financial markets is even less certain. This uncertainty is compounded by the fact that forward-looking markets have already priced in high hopes for this technology. If productivity and earnings gains fail to materialize on the expected scale or timeline, a meaningful repricing could follow.

¹This note builds on our colleagues' earlier work in '[Macro, Markets, and Machines](#)' where they provide a thorough review of existing literature and enterprise-wide insights outlining AI's potential impact on global markets and public policy through an economic and investment lens. The report highlights AI's promise for productivity and profits, while also emphasizing regional differences, the risks of labour displacement, supply bottlenecks, policy challenges, and valuation concentration.

²Computers and peripheral equipment and data-centre structures are used as proxies for AI-adjacent capital spending in the U.S. They are broad categories, however, which cover more than AI-specific investments.

In such an environment, **scenario analysis is a useful framework**. Rather than relying on a single forecast, it allows us to explore a range of plausible outcomes, assess risks around the baseline, and see if we can draw general conclusions. In this note, we explore three broad scenarios through our Scotia macro model: one where productivity gains are significant, asset valuations continue to grow, and impact on labour markets is neutral; one where productivity gains are significant but labour markets face disruptions and displacement; and one where expected gains fail to materialize and equity markets reprice sharply. Before moving to these scenarios, we discuss how AI is likely to impact the economy.

AI AND PRODUCTIVITY: THE TASK-COST CHANNEL

At its core, AI operates through familiar macroeconomic channels. Like previous technological advances, it acts primarily as a productivity shock, with spillovers to labour markets, asset prices, capital spending, and energy demand. What differentiates AI is the uncertainty around the balance of these effects and the speed at which they unfold.

AI raises productivity by lowering the cost of completing certain tasks. In a [task-based framework](#), the aggregate gain depends on the proportion of tasks impacted and the average improvement (or task-cost saving). For example, if AI meaningfully improves 10% of tasks in a firm with a 50% efficiency gain, the productivity of that firm would increase by 5%, all else equal. At the macro level, however, the gains are filtered through adoption frictions, firm capabilities, worker training, and regulations.

The impact of AI on the labour market is highly uncertain, with the most material implications for the transmission of productivity gains through the economy. Ultimately, the technology's impact on the share of labour in the economy will come down to whether it automates (substitutes) or augments (complements) certain tasks. Automation reduces labour demand in affected areas and lowers the labour share of income, while augmentation enhances the productivity of workers who remain involved in affected tasks. As with past technological advances, automation might be balanced in the long-run by the creation of new tasks (reinstatement) that can be considered as new AI-related tasks, such as tool design, model oversight, or maintenance. However, this offset is difficult to measure as it is highly uncertain whether new task creation occurs quickly enough, and at sufficient scale, to offset displacement in lower-complementarity roles. Rapid adoption would cause short-term stress in labour markets as workforce displacement happens faster than job creation and retraining. Lower wages could help absorb displaced workers, but since wage adjustment is slow, fast adoption would temporarily increase unemployment, perhaps sharply. Conversely, gradual implementation would allow for retraining and natural attrition would soften the impact.

For equities, the mechanism is clear: if AI raises anticipated growth rates or profit margins, stock valuations remain high; otherwise, they're at risk given currently heightened and concentrated valuations. Current market prices reflect assumptions about how far and long productivity improvements will last. We [estimate](#) that about 20% of the US stock market value is related to AI expectations, reflecting these high hopes for future growth. If the technology proves better than current market expectations (and bolsters profits expectations), asset valuation would continue to rise. In contrast, failure to deliver at the expected scale would lead to a sharp repricing.

Token costs introduce another uncertainty around the transmission of AI to productivity, margins, and valuations. If scaled AI use proves more expensive than firms expect, adoption may either slow or shift toward lower-cost models, compressing provider margins and increasing the risk of an equity repricing. High usage costs could also reveal adoption inefficiencies if firms are unable to translate AI spending into measurable productivity gains, with similar risks to valuations. Conversely, if productivity gains are large enough to justify the higher operating costs, firms may have a stronger incentive to substitute capital for labour, reinforcing the labour-displacement channel explored in the upside scenario with labour market disruption.

There is also the question of the distributional effects of AI-adoption as income shifts from labour to capital, with certain tasks and income quintiles more exposed than others. If more lower income occupations fall into the high-exposure low-complementarity segment than high income occupations, clear welfare and distributional implications arise. Indeed, while Statistics Canada [show](#) that exposure to AI increases with employment income, higher earners hold jobs that have more complementarity (smaller automation risk) than lower- and mid-income jobs. In the final section of [Macro, Markets, and Machines](#), Rebekah Young dives into the welfare considerations of AI adoption, and offers a policy playbook on smoothing the labour transition and facilitating and incentivizing complementarity- and diffusion-focused adoption.

SCENARIO ANALYSIS

Given the wide dispersion of possible outcomes, scenario analysis provides a useful lens to assess risks and see if general conclusions can be drawn. Calibrating these scenarios is extremely tricky and imprecise, and the numerical results should be interpreted with caution. Still, scenarios are informative in illustrating directional macroeconomic impacts and the importance of different channels. The value of this

exercise then is not in delivering an in-house estimate on the impact of AI. Rather, the goal is to provide an outline of the frameworks and mechanisms that will inform our thinking around AI-adoption and how it might impact our forecasts over time.

Charts 2–9 below summarize the resulting deviations from baseline across the three scenarios for key U.S. and Canadian macro variables.

Chart 2

U.S. GDP




Chart 3

PCE Inflation


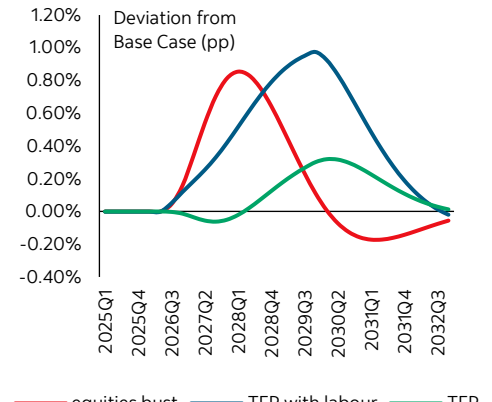


Chart 4

U.S. Unemployment



Source: Scotiabank Economics.

Source: Scotiabank Economics.

Source: Scotiabank Economics.

Chart 5

Federal Fund Rate




Chart 6

Canada GDP


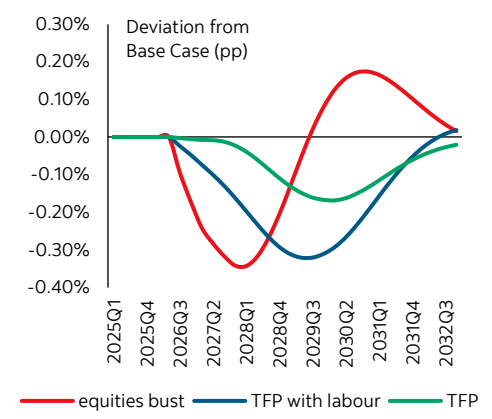


Chart 7

CPI Inflation



Source: Scotiabank Economics.

Source: Scotiabank Economics.

Source: Scotiabank Economics.

Chart 8

Canada Unemployment





Chart 9

Overnight Rate



Source: Scotiabank Economics.

Source: Scotiabank Economics.

UPSIDE SCENARIOS: PRODUCTIVITY ACCELERATION WITH DIFFERING LABOUR MARKET OUTCOMES

On the upside, AI could deliver **stronger-than-expected productivity gains**, particularly if adoption proves more complementary to existing capital and the labour force than currently assumed. Equity markets would also continue their upward trajectory, supported by higher expected earnings rather than purely valuation effects.

The labour market implications however hinge critically on the degree of complementarity between labour and AI, as well as the speed of adoption and creation of new tasks. Given the high uncertainty of this part of the AI-transmission channel, and given the importance of the unemployment assumption for the simulation results, we separate the upside scenario into two parts.

SCENARIO 1A: LABOUR-NEUTRAL PRODUCTIVITY ACCELERATION

Scenario 1a assumes AI raises productivity without disrupting labour markets. In other words, labour augmentation and creation of new tasks are assumed to offset the effects of automation without materially disrupting the level of employment. Or it could be that displacement that does occur is largely absorbed by demographic forces—namely, AI substitutes for retiring workers rather than actively displacing the existing workforce. We view this as a relatively low-probability outcome, but it is useful for illustrating the mechanics of a labour-neutral productivity shock before layering in labour-market disruption in Scenario 1b.

The model simulations are run for both the U.S. and Canada, with the U.S. used as the benchmark economy given its central role in AI development, investment, and existing literature. We calibrate the benchmark U.S. TFP growth gain at approximately +0.5 percentage points per year, sustained over five years. Canada's shock is scaled down to +0.3 percentage points per year over the same time horizon, reflecting differences in AI-adoption metrics (for details on calibration of this shock, please see Box 1).

In this case, the macroeconomic impulse is driven largely by two channels. First, higher TFP raises the economy's long-run speed limit by lifting potential output and the long-run level of investment. In isolation, this is a positive supply shock: capacity expands faster than demand, creating excess supply and near-term disinflation. Demand ultimately catches up. Higher income boosts consumption, while AI-driven innovation creates new goods and services that generate incremental demand. Second, a temporary investment impulse that captures the AI-build phase. This reflects stronger-near term demand for capital spending, data centres, electricity capacity, specialized labour, and other inputs. As a result, investment initially overshoots its long-run trend before slowing in the medium term and eventually converging to a higher level consistent with stronger productivity growth.

This initial investment impulse matters for the inflation profile. Rather than producing immediate disinflation, the initial build-out phase adds demand at the same time that supply capacity is rising. The result is a roughly flat inflation response initially, before disinflationary pressure emerges as the productivity channel dominates. This short-run demand and inflationary offset is likely more acute in the U.S. than in Canada, given the significantly greater scale of AI-related capital spending and infrastructure build-out in the U.S.

Wages also rise in this scenario, consistent with the complementarity between AI and labour. However, these gains largely reflect stronger productivity rather than an increase in labour costs. In other words, workers earn more because they produce more, limiting the inflationary impact.

The net policy rate response in this scenario is also limited. In the U.S., the policy rate response is broadly flat, as the eventual disinflationary impulse from stronger productivity is largely offset by a higher neutral rate.³ In Canada, the model implies a 25 bp cut, but it is not long lasting.

Overall, U.S. GDP is lifted by approximately 3% by 2032. Canadian GDP rises by approximately 1.5% over the same period, with the smaller gain reflecting slower and narrower AI-adoption in Canada, owing to lower readiness and historically weaker TFP growth.

SCENARIO 1B: TFP INCREASE WITH LABOUR DISRUPTION

As in prior periods of technological innovation, labour market disruptions are expected as adoption likely occurs at a more rapid pace than reinstatement. The key uncertainty is the magnitude of the disruption. Therefore, Scenario 1b should be interpreted primarily as illustrating the direction and relative sensitivity of the macroeconomic effects to the labour market adjustment, rather than as a precise estimate of the expected increase in AI-driven unemployment.

³ We allow our model's neutral rates to rise endogenously with potential output growth, using the [Laubach-Williams](#) method, with a coefficient of one between potential growth and the neutral rate. The neutral rate rises by roughly 50 bps in the U.S. and 25 bps in Canada, partly offsetting the easing impulse implied by our model in the medium-run in response to the disinflationary effects of higher productivity.

This scenario adds an AI-related rise in the unemployment rate on the order of 0.3 to 0.4 percentage points in the U.S. and Canada, with the U.S. impact at the upper end of the range.⁴ This reflects the same broad exposure risk, but with wider and faster expected adoption and build out in the U.S., which could translate into somewhat greater worker displacement.

Even in this more disruptive scenario, the net impact on economic activity would likely remain positive relative to the baseline. However, the gain materializes more slowly than in Scenario 1a as layoffs weigh on household incomes and spending, partially offsetting the positive effects of the initial investment impulse and stronger productivity. This moves the economy more clearly into excess supply and, combined with easing cost pressures directly from AI adoption, generates a stronger disinflationary impulse, pointing to lower interest rates than in Scenario 1a.

The model implies a larger monetary policy easing response in the U.S. compared to Canada, reflecting a larger TFP and unemployment shock, both adding to excess supply. While the initial investment impulse is also larger, it does not fully offset the additional supply capacity from stronger productivity growth.

Here, the U.S. GDP gain is roughly half as large as in Scenario 1a in the short to medium run. Over time, the level of GDP gradually catches up as productivity gains continue to accumulate, leaving U.S. GDP around 3% above baseline by 2032. On the other hand, Canada's GDP initially falls about 0.3% below baseline, as labour markets disruptions more than offsets the positive productivity gains, before gains accumulate and GDP eventually rises roughly 1.5% above baseline over the same period.

DOWNSIDE SCENARIO: VALUATION CORRECTION AND CONFIDENCE SHOCK

The key downside risk is that AI fails to deliver the expected productivity gains or sustained profit growth within a reasonable timeline. In that scenario, the recent AI-driven equity rally would likely unwind over a relatively short horizon—potentially within a year—leading to a broad tightening in financial conditions. While there would be a modest drag on TFP as anticipated efficiency gains fail to materialize, the primary transmission channel to the macroeconomy would operate through financial markets.

We calibrate this in our models as a U.S.-led repricing of equities of roughly 20%, driving a somewhat smaller correction in Canadian equities of around 15%, alongside a temporary increase in risky spreads, and a confidence shock that directly weighs on consumption and investment. We also include a negative TFP growth shock to capture the unwinding of currently embedded expectations of stronger future productivity.

In particular, a repricing of equities would raise the cost of capital and weaken balance sheets, especially in sectors most exposed to the AI investment cycle. Investment would likely slow materially, with an outsized impact on AI-related capital spending that has been supported by elevated valuations and optimistic earnings expectations.

The tightening in financial conditions would also spill over to households. Higher borrowing costs and negative wealth effects from equity market declines would weigh on sentiment and consumption, reinforcing the slowdown in activity.

Taken together, this dynamic resembles a conventional negative demand shock. Softer domestic demand would ease inflationary pressures, while the deterioration in growth momentum would shift the policy bias toward easing. As a result, both inflation and policy rates would likely move lower relative to baseline.

In this scenario, U.S. GDP falls by roughly 2% below baseline by the end of 2027, while Canada's GDP is around 1.75% lower. The somewhat larger U.S. impact reflects its greater exposure to AI-related equity valuations and investments, with spillover effects onto Canada.

PUTTING IT TOGETHER

The macroeconomic implications of AI ultimately hinge on two key questions: whether productivity gains materialize, and how quickly labour markets and demand adjust to the technology and additional supply.

If AI delivers on current expectations, the economy would benefit from stronger productivity growth but the adjustment process may not be smooth. If, instead, AI fails to generate meaningful efficiency gains, the dominant channel becomes financial and the resulting slowdown would resemble a conventional demand shock.

Taken together, a consistent conclusion emerges: AI is likely to be disinflationary across a wide range of outcomes.

⁴ We calibrate the unemployment shock using a forthcoming Bank of Canada estimate [suggesting](#) that the unemployment rate for AI-exposed occupations could rise by around 1.2 percentage points. Since roughly 30% of workers in both economies are in high-exposure, low-complementarity occupations, this implies an aggregate unemployment rate shock of roughly 0.3 to 0.4 percentage points.

Box 1: Calibration of AI TFP Shocks for the U.S. and Canada

For calibration purposes, we use the U.S. as our benchmark as it is the focus of most large existing studies. The range of macro estimates from existing studies is wide, with [Acemoglu \(2024\)](#) at the lower end of that range with an estimated total factor productivity (TFP) gain of 0.7% to 1% within the next ten years (or 0.07%–0.1% annually). More optimistic estimates point to much larger effects. Conclusions are heavily impacted by assumptions, and vary [even](#) when using similar conceptual framework.

For the upside scenarios of improved productivity, we use a mid-range assumption of a U.S. TFP shock of roughly 0.5 percentage points per year over five years. This assumption is based on our colleagues [survey](#) of available estimates where they identify 0.5% as the average annual TFP-growth boost across the studies reviewed.

For Canada, we scale the U.S. shock using the exposure, preparedness, and access framework defined by the [IMF](#), while also applying a further penalty for Canada's historically weaker productivity performances. In the IMF framework, exposure refers to the share of workers and industries affected by AI-adoption, regardless of whether through automation or augmentation; preparedness refers to the readiness of institutions, structures, and workforce; and access refers to the availability of the needed technology and infrastructure.

A 2021 census-based [study](#) by Statistics Canada suggests that 31% of Canadian workers are in high-exposure low-complementarity occupations, and 29% are in high-exposure high-complementarity occupations. This puts the share of highly exposed workers at 60%, broadly in line with [estimates](#) of labour exposure in advanced economies and the U.S. However, similar exposure does not equal similar productivity gains.

Canada ranks below the U.S. on the IMF [AI Preparedness Index](#) (0.71 versus 0.77), has materially [lower](#) AI-related investment and domestic infrastructure build-out, and has [historically](#) lagged the U.S. in TFP growth, including during the period of high investment in information and communication technologies. Assuming comparable access, as Canadian firms can import much of the compute needed for AI adoption from the U.S., but a weaker domestic scaling channel, we apply a conservative Canadian TFP shock of roughly 0.3 percentage points per year over five years.

This scaling also reflects recent policy efforts to strengthen Canada's domestic AI capacity, including the federal government's [national AI strategy](#) and the [Canadian Sovereign AI Compute Strategy](#), which aim to mobilize infrastructure investment and expand domestic compute capacity. As a result, we apply a smaller penalty than Canada's historical productivity.

This report has been prepared by Scotiabank Economics as a resource for the clients of Scotiabank. Opinions, estimates and projections contained herein are our own as of the date hereof and are subject to change without notice. The information and opinions contained herein have been compiled or arrived at from sources believed reliable but no representation or warranty, express or implied, is made as to their accuracy or completeness. Neither Scotiabank nor any of its officers, directors, partners, employees or affiliates accepts any liability whatsoever for any direct or consequential loss arising from any use of this report or its contents.

These reports are provided to you for informational purposes only. This report is not, and is not constructed as, an offer to sell or solicitation of any offer to buy any financial instrument, nor shall this report be construed as an opinion as to whether you should enter into any swap or trading strategy involving a swap or any other transaction. The information contained in this report is not intended to be, and does not constitute, a recommendation of a swap or trading strategy involving a swap within the meaning of U.S. Commodity Futures Trading Commission Regulation 23.434 and Appendix A thereto. This material is not intended to be individually tailored to your needs or characteristics and should not be viewed as a “call to action” or suggestion that you enter into a swap or trading strategy involving a swap or any other transaction. Scotiabank may engage in transactions in a manner inconsistent with the views discussed this report and may have positions, or be in the process of acquiring or disposing of positions, referred to in this report.

Scotiabank, its affiliates and any of their respective officers, directors and employees may from time to time take positions in currencies, act as managers, co-managers or underwriters of a public offering or act as principals or agents, deal in, own or act as market makers or advisors, brokers or commercial and/or investment bankers in relation to securities or related derivatives. As a result of these actions, Scotiabank may receive remuneration. All Scotiabank products and services are subject to the terms of applicable agreements and local regulations. Officers, directors and employees of Scotiabank and its affiliates may serve as directors of corporations.

Any securities discussed in this report may not be suitable for all investors. Scotiabank recommends that investors independently evaluate any issuer and security discussed in this report, and consult with any advisors they deem necessary prior to making any investment.

This report and all information, opinions and conclusions contained in it are protected by copyright. This information may not be reproduced without the prior express written consent of Scotiabank.

™ Trademark of The Bank of Nova Scotia. Used under license, where applicable.

Scotiabank, together with “Global Banking and Markets”, is a marketing name for the global corporate and investment banking and capital markets businesses of The Bank of Nova Scotia and certain of its affiliates in the countries where they operate, including: Scotiabank Europe plc; Scotiabank (Ireland) Designated Activity Company; Scotiabank Inverlat S.A., Institución de Banca Múltiple, Grupo Financiero Scotiabank Inverlat, Scotia Inverlat Casa de Bolsa, S.A. de C.V., Grupo Financiero Scotiabank Inverlat, Scotia Inverlat Derivados S.A. de C.V. – all members of the Scotiabank group and authorized users of the Scotiabank mark. The Bank of Nova Scotia is incorporated in Canada with limited liability and is authorised and regulated by the Office of the Superintendent of Financial Institutions Canada. The Bank of Nova Scotia is authorized by the UK Prudential Regulation Authority and is subject to regulation by the UK Financial Conduct Authority and limited regulation by the UK Prudential Regulation Authority. Details about the extent of The Bank of Nova Scotia's regulation by the UK Prudential Regulation Authority are available from us on request. Scotiabank Europe plc is authorized by the UK Prudential Regulation Authority and regulated by the UK Financial Conduct Authority and the UK Prudential Regulation Authority.

Scotiabank Inverlat, S.A., Scotia Inverlat Casa de Bolsa, S.A. de C.V., Grupo Financiero Scotiabank Inverlat, and Scotia Inverlat Derivados, S.A. de C.V., are each authorized and regulated by the Mexican financial authorities.

Not all products and services are offered in all jurisdictions. Services described are available in jurisdictions where permitted by law.