AI's Impact on the Workforce — PHU01

Executive Summary (condensed, English)

How agentic AI could affect the need for labour in Sweden through 2035

Per Högberg, Chairman of the Municipal Executive Board, Mullsjö Municipality, Sweden

PHU01 · Condensed from the full Swedish-language report, version 2.2 (UTGIVEN / published), 19 November 2026

About this summary

This is a condensed executive summary of a full Swedish-language investigation (approx. 80 pages, 13 chapters plus appendices). It is prepared for an international, non-Swedish-speaking audience and omits the full evidence tables, source register and sensitivity appendices of the original. Every figure below is taken directly from the full report; none has been recalculated or rounded differently. A complete English translation of the full report is available on request.

Who wrote this, and how

The investigation was produced by Per Högberg, chairman of the municipal executive board (kommunstyrelsens ordförande) of Mullsjö Municipality, Sweden, on his own initiative between June and November 2026. It was not commissioned by any authority or organisation. It was produced in collaboration with Claude models from Anthropic (primarily Claude Sonnet 5.5, with review by Claude Opus 5.5 and Fable 5.2, and some review support from ChatGPT): the AI searched and compiled sources, built the calculation model to the author's instructions, and drafted chapter text, but every assumption in the model and every conclusion is the author's own. Anthropic is itself a party to the development the report describes; no data from Anthropic is used as supporting evidence.

Per Högberg's background combines Swedish municipal governance (chairman of the executive board, Mullsjö, 2023–2026) with 2007–2023 in the IT sector focused on business-process development, and engagement with the SAP user community since 1998 — including founding and leading SUGEN (SAP User Group Executive Network) — plus work as an industry researcher in information modelling at Jönköping University since 2006.

Why this exists — and what it is not

Agentic AI has only begun to see wider use during 2026. Measured effects to date, in Sweden and internationally, concern almost exclusively generative AI used as a chat interface — they show no broad transformation of employment yet, though a weakened entry path into AI-exposed occupations for young people is visible. There is therefore no evidence base to answer the central question directly. This is why the investigation is explicitly a scenario study, not a forecast: the effect levels are the author's own assumptions, stated and justified openly, not measurements. The author recommends the calculation be updated roughly every six months as real-world outcomes become measurable.

AI's greatest power is judged to be analysis, not administration: agentic AI has most impact where large volumes of data can be gathered and analysed to surface patterns people often miss — in casework and expert work, but also in care, where continuously collected data could flag early signs that need action. Documentation and administrative support are the easiest effects to introduce first, but are assessed as the least important effect (see chapter 3.2 of the full report).

Headline results

Sweden's Public Employment Service (Arbetsförmedlingen, AF) projects an additional labour shortfall of roughly 270,000 people by 2035 (270,972 in AF's own base-scenario table) — the gap between the labour demand AF projects and the number of people who would be employed if nothing else changes. The investigation asks how much of that gap agentic AI could help close, by letting the same work be done by fewer people.

- To close the entire labour shortfall, average labour demand would need to fall by just 5.6% compared with today's employment — a figure that follows directly from AF's and Statistics Sweden's (SCB) own data and involves no assumption about AI at all.
- The three scenarios modelled give a gross reduction in labour demand of 11.0%, 17.9% and 24.6% (Low, Mid, High) — roughly 532,000, 866,000 and 1,192,000 people. All three are well above the 5.6% break-even.
- Without people moving between occupations, the shortfall is covered to 82% (Low) and 99% (Mid) when calculated across 46 occupational groups; around 48,000 people of shortfall remain in Low, mainly in care, health care, construction and services. On AF's own finer 147-group level, coverage is 73% (Low) and 95% (Mid), leaving about 74,000 and 15,000 people of shortfall respectively. In care's largest shortfall occupations the break-even requirement is 14–16%; the Low-scenario assumption of 10% is not sufficient on its own, while the Mid-scenario assumption of 20% is — a result that holds at both levels of occupational detail and is treated as the robust finding.
- Beyond closing the shortfall, the scenarios free up additional labour capacity equivalent to roughly 261,000, 595,000 and 921,000 employees (Low, Mid, High) — a surplus. In High this equals 19% of today's workforce and 8.7% of Sweden's population. This surplus arises mainly in office, casework and expert occupations, where the shortfall is small. These figures are not an unemployment forecast: how much of the freed capacity turns into fewer people employed depends on occupational mobility, new demand, new tasks and working hours.
- If lower prices from efficiency gains stimulate extra demand ("volume compensation"), the effect shrinks. The main calculation assumes no volume compensation and is therefore an upper bound. How large the real compensation will be cannot be determined with today's evidence; even at 50% compensation, 89% of the Mid-scenario shortfall is still covered.

The central finding

The question is not whether AI can free up enough labour capacity in aggregate — it can, on every scenario modelled. The question is whether that capacity is freed where the shortfall actually is. Sweden could face a labour shortage and a labour surplus at the same time: the shortfall sits mainly in care and health care (37.5% of the entire shortfall), while freed capacity arises mainly in casework, office and expert occupations. This makes three things the central political questions: enabling occupational mobility, directing AI specifically at care and health care, and — if a real surplus materialises — how to manage it (working hours, retirement age, financing and the tax base).

Key caveats the investigation is explicit about

- Every figure describing something that has already happened is classified by evidence level: level 1 = supply/capacity (a vendor offers something), level 2 = adoption (something is actually used,

measured), level 3 = outcome (a measurable effect). A vendor's own claim about its own product is always level 1, and the report never presents such a claim as level 2 or 3 evidence.

- Volume-compensation demand effects, demand shifts elsewhere in the economy, and entirely new tasks and occupations are acknowledged as real compensating mechanisms but are only partly or not at all quantified, for lack of AI-specific evidence.
- Public-sector activities with a large public share — for example defence, and parts of IT and expert services — are modelled as market-driven, on the simplifying assumption that the public sector largely follows the same efficiency pattern as the private sector, possibly with some lag. This simplification is explicitly flagged as not separately verified and as something a future revision may need to test.
- The report's own introduction is explicit that it cannot say how fast this unfolds, how much of the freed capacity will be absorbed by new demand and new tasks, and whether the required occupational mobility is achievable in practice. It is presented as a first quantitative basis to update, not a final answer.

What this means for policy — in order

- First: direct AI specifically at the shortfall occupations, above all care and health care, and measure the effect.
- Second: make occupational mobility possible at a scale today's systems are not built for, and make use of people currently outside the workforce.
- Only if the freed capacity is still not absorbed do working hours, retirement age and the tax base become the central questions — these need preparation now, because they take a long time to work through, but decisions on them should be conditional on what the ongoing monitoring actually shows.

About the full report

The full report (in Swedish) runs to roughly 80 pages across 13 chapters plus six appendices: occupation-type classification and assumptions (A), results by occupational group at two levels of granularity (B), the reasoning behind every assumption by impact type (C), a full source register (D), a glossary (E), and the author's earlier conceptual background material on municipal agentic AI (F). All quantitative claims are traceable to a named, logged source. A companion investigation, PHU02, on the development of elderly care, is planned as a follow-up.

A complete, unabridged English translation of the full report can be produced on request.