A Report on "Gender-Biased Technological Change: Milking Machines and the Exodus of Women from Farming" by Ager et al. (2026)

Reviewer 2

February 11, 2026

v1

isitcredible.com

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made here. Any errors in the report will be corrected in future revisions.

> I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.
>
> Plato, *The Apology of Socrates*, 21d

> To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.
>
> Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Citation: Ager, P., Goñi, M., and Salvanes, K. G. (2026). Gender-Biased Technological Change: Milking Machines and the Exodus of Women from Farming. *American Economic Review*. Vol. 116, No. 1, pp. 246–286.

URL: https://doi.org/10.1257/aer.20240167

Abstract Summary: This paper examines how the widespread adoption of milking machines in twentieth-century Norway, a gender-biased technological change, affected women's work and economic status. The study finds that the machines displaced young rural women from farming, leading them to migrate to cities, acquire more education, and find better-paying, skilled employment, which ultimately improved their economic status relative to men.

Key Methodology: Instrumental Variables (IV) approach using a shift-share design, exploiting plausibly exogenous variation in local uptake of milking machines by combining nationwide adoption with local dairy farming intensity in 1930, applied to individual-level Norwegian registry data.

Research Question: How did the automation of farming tasks, specifically the adoption of milking machines, contribute to the transformation of women's work in the twentieth century, and did it bring economic hardship or gains to affected women by pushing them out of agriculture?

Summary

Is It Credible?

Ager et al. present a compelling argument that the introduction of milking machines in Norway acted as a gender-biased technological shock that, paradoxically, improved women's long-term economic status. By automating a task traditionally performed by women, the technology "pushed rural young women in dairy-intensive areas out of farming" and into urban centers, where they accessed education and better-paying jobs (p. 246). The central claim—that displacement from a traditional sector led to long-term income gains for women relative to men—is supported by a rigorous instrumental variables strategy. The magnitude of the effects, including a narrowing of the gender income gap by about 2 percentile ranks and a reduction in the labor force participation gap by more than 3 percentage points, suggests a substantial structural shift rather than a marginal adjustment (p. 272).

The credibility of the main causal claim is high, though it is bounded by specific contextual factors. The authors candidly note that the positive long-term outcomes were "only possible because Norway provided the local schooling infrastructure" and a booming public sector to absorb the displaced labor (p. 281). This indicates that the article provides credible evidence for what can happen under specific institutional conditions—strong safety nets and expanding public employment—rather than a universal law of automation. The reliance on income measurements at different ages (45, 52, and 62) for different cohorts introduces some noise into the long-term estimates. However, the authors' use of percentile ranks—which rank individuals within their own birth cohort—effectively adjusts for life-cycle effects and inflation, mitigating concerns that this drives the results (p. 256). Furthermore, while the authors use historical census data to demonstrate a lack of pre-trends, the shift from high-quality modern registry data to lower-quality historical data for these placebo

tests introduces a degree of uncertainty, as noise in the historical record could potentially mask underlying trends (p. 261). It is worth noting, however, that the authors also provide an event-study analysis using high-quality registry data for later cohorts, which shows parallel trends prior to the shock (p. 263).

The evidence regarding the mechanisms of change is mixed. The claim of "occupational upgrading" is complicated by selection bias; the analysis showing women moved into higher-skilled jobs is conditional on them leaving agriculture, a decision that is itself an outcome of the treatment (p. 275). Women who left farming might have been the most capable or ambitious regardless of the machine. However, the authors bolster this with an analysis of educational attainment on the full sample, finding a 1.5 percentage point increase in undergraduate education (p. 276). This unconditional result is methodologically sound and provides the necessary support for the human capital mechanism. Conversely, the demographic mechanism—that automation reduced fertility—appears less central. The estimated effects are "modest," and data limitations prevent the researchers from observing the spouse's exposure to the shock, making this a less robust component of the overall narrative (p. 278).

The Bottom Line

The article provides credible evidence that gender-biased automation in mid-20th century Norway displaced women from agriculture but ultimately improved their economic standing relative to men. The findings regarding migration and educational investment are robust, though the specific claims about occupational upgrading are slightly weakened by selection issues. Readers should recognize that these positive welfare gains were heavily dependent on the specific institutional context of Norway's expanding public sector and education system.

Potential Issues

Selection bias in the occupational upgrading analysis: The article argues that a key mechanism for women's long-term gains was occupational upgrading into higher-skilled jobs. However, the analysis supporting the occupational upgrading claim is restricted to "individuals who did not work in agriculture as middle-aged adults" (p. 275). Since the decision to leave agriculture is a primary outcome of the technological shock, conditioning the sample on this choice may introduce selection bias. The analysis therefore estimates the effect of the shock on the selective group of women who left farming, who may differ systematically in unobservable characteristics like ambition or ability from those who remained. The authors appear to recognize this limitation by providing a complementary analysis of educational attainment. As they note, "individuals who remained in agriculture could, in principle, also invest more in their education," and they find that exposure to the shock increased the likelihood of obtaining an undergraduate degree in an analysis of the "full baseline sample" (pp. 275–276, Table 4). This second finding supports the mechanism without conditioning on the decision to leave farming.

Limited generalizability due to institutional context: The article's conclusion that gender-biased technological change can resolve labor misallocation and improve women's economic status is heavily dependent on the unique historical context of mid-20th century Norway. The positive long-term outcomes for displaced women were contingent on the simultaneous expansion of the Norwegian public sector, which created a high demand for educated female labor in new white-collar occupations like teaching and nursing (p. 274). The authors explicitly identify this as a key scope condition for their findings, stating that the positive outcome "was only possible because Norway provided the local schooling infrastructure for women to invest in education and take on new and better job opportunities" and concluding that "the effects of automation are institution dependent" (p. 281). This

suggests the findings may not generalize to settings without a similar, concurrent expansion of female-intensive service or public sector employment, where displaced agricultural workers might have faced worse outcomes.

Debatable validity of the placebo test due to data quality: The article's identification strategy relies on the absence of long-term pre-existing trends in outcomes between dairy-intensive and non-dairy-intensive regions. To support this, the authors conduct a placebo test using historical census data from 1900 and 1910, finding no significant pre-trends (pp. 261–262). The main analysis, however, uses high-quality modern registry data with "unique personal identifiers," which the authors note is superior to the "automated linking methods" required for the historical census data (p. 255). While the authors are transparent about the different data sources, they do not discuss whether the lower quality of the historical data used for the placebo test—which is more susceptible to measurement error—could bias the placebo coefficients toward zero and thus mask the existence of genuine pre-trends. This unacknowledged difference in data quality between the main analysis and this crucial validation test introduces some uncertainty about the strength of the evidence against long-term confounding trends.

Measurement of the primary long-term outcome: The study's main long-term outcome is income percentile rank, but due to data availability, it is measured at different ages for different cohorts: age 45 for the 1950–1970 cohorts, age 52 for the 1940 cohort, and age 62 for the 1930 cohort (p. 256). This introduces a potential measurement issue, as income ranks may not be perfectly comparable across different points in the life cycle, particularly for women whose careers may be non-linear. The authors acknowledge this issue and provide three pieces of evidence to mitigate concerns: they argue that "percentile ranks are less sensitive to the age at which income is measured than the income in levels," they show that income ranks at these different ages are highly correlated in their data, and they demonstrate that their results are robust to excluding the older cohorts for whom income is measured at later ages (pp. 256–257,

Supplemental Appendix Figure A.6, Supplemental Appendix Table A.6).

Potential omitted variable bias in fertility analysis: The article posits that the technology shock reduced fertility by increasing women's opportunity cost of having children. The analysis models fertility, a joint household decision, based only on the woman's childhood exposure to the shock, omitting the characteristics of her spouse (p. 278). This could lead to omitted variable bias if women who migrated from dairy-intensive areas systematically married men with different attributes, such as more educated or urban-born men, who had different fertility preferences. The authors acknowledge this as a data limitation, stating that "the structure of the registry data does not allow us to measure the exposure to milking machines of wife and husband separately" (p. 278). While changes in spousal characteristics may be a mechanism rather than a confounder for the *total* effect on the woman's life, the inability to disentangle these factors limits the precision of the fertility mechanism specifically.

Interpretation of the fertility mechanism's importance: The article presents changes in fertility and marriage patterns as a "second complementary mechanism" facilitating women's labor reallocation (p. 277). However, the estimated effects are quantitatively small. A one standard deviation increase in milking machine exposure is associated with a fertility reduction of about 3 percent of the sample mean and an increase in the age at first birth of about 1 percent. The authors are transparent about this, stating that "the effect is modest" and "quantitatively small" (p. 278). While statistically significant and consistent with the broader narrative, the practical importance of these effects as a major driver of the observed economic transformation appears limited.

Minor clerical and presentational issues: There are several minor presentational inconsistencies in the text. First, the abstract states that "differences in LFP rates dropped by almost 4 percentage points" (p. 248). While technically accurate as a rounding of the 3.2 percentage point differential effect found in Table 3 (0.038 for women minus 0.006 for men), the abstract's summary is slightly less precise than

the body text (p. 272). Second, there is a small discrepancy in the reported number of observations for the short-term income analysis, cited as 450,350 in the main text (p. 265) and in a supplemental figure (p. 48), but as 450,325 in the notes to Figure 4 and Table 1 (pp. 267, 268). These minor issues do not affect the substantive conclusions of the article.

Future Research

Testing institutional dependencies: Future work should apply similar identification strategies to contexts where the public sector or educational infrastructure was not expanding simultaneously with agricultural mechanization. Comparing the Norwegian experience to a context with a stagnant service sector would clarify whether automation yields welfare gains for displaced women in the absence of strong institutional support.

Bounding selection bias in occupational shifts: To address the selection bias inherent in analyzing occupational upgrading conditional on leaving agriculture, researchers could employ bounding exercises (e.g., Lee bounds). This would provide a range of estimates for the effect of automation on job quality that accounts for the non-random selection of workers who exit the farming sector.

Household-level mechanisms: Research utilizing data that links spouses prior to marriage or observes full household histories could better test the fertility and marriage market mechanisms. This would allow for a more precise estimation of how gender-specific labor shocks affect household formation and bargaining power, avoiding the omitted variable bias present when only one partner's exposure is observed.

isitcredible.com