AI Agents and Higher-Order Work

Suproteem K. Sarkar
University of Chicago

Please click here for the latest version

How do AI agents influence knowledge work? This paper finds that agents shift worker effort from implementation to supervision, which especially benefits verifiable work and expert workers. I analyze data from the coding platform Cursor to study agents in software production. First, I find that workers have become less likely to produce output manually after AI agents were introduced, and are more likely to delegate work to agents. Second, workers use agents for abstract, higher-order tasks like delegation, context gathering, and planning. Third, agents are used more frequently in settings with easier-to-verify work outputs. Fourth, experienced workers appear more skilled at delegating work to agents—they ask fewer questions, seek more alignment through planning, and accept agent outputs at higher rates. Fifth, using variation from the platform's feature release timeline, I find that agents increase software output, especially for firms with more verifiable work and more experienced workers. The complementarity between AI agents and expertise contrasts with evidence on non-agentic AI tools, which workers are more likely to use for support than for delegation. Taken together, the results suggest that the returns to expertise may rise as knowledge work becomes more abstract.