

Scaling AI-assisted code review with practical defaults

Paul Gendek | Senior Staff Engineer

At The Estée Lauder Companies, I initially enabled GitHub Copilot code review across more than 30 repositories, including front-end services, component libraries and shared platform tooling.

As we used it, a tradeoff emerged: repeated automatic reviews during development were consuming developers' AI usage allowances. Review frequency needed to reflect how each repository was used.

We adjusted the defaults in a small number of repositories. In some, we disabled automatic reviews on draft pull requests. In others, we made reviews opt-in, allowing a reviewer to request a final sweep once the PR was ready.

My contribution included enabling broad coverage and helping refine the workflow around its practical constraints. We retained broad access to AI-assisted feedback while applying more selective defaults where repeated reviews were consuming too much of the available allowance.

The lesson was about engineering judgment: useful adoption requires attention to timing, resource limits and team workflows. Enabling a capability creates the starting point; observing its use helps determine the right defaults.