Real Productivity, Rare Pushbacks: 5 Principles for AI Agents in Product Development Teams

Vivek Ganesan Kamal Raj Sekar
Ampyard Private Limited
{vivek.ganesan, kamalraj.sekar}@ampyard.com

ABSTRACT

The integration of AI agents into team workflows holds immense potential for boosting productivity, yet current practices often lead to suboptimal outcomes and even resistance. Our analysis of AI agent deployment reveals several shortcomings: a lack of focus on bottleneck-focused optimization, neglect of unpleasant or high-value tasks, accumulating work to be done towards the end of the release cycle, and insufficient human oversight, all of which hinder efficiency and team acceptance. We propose a strategic framework anchored in five principles to address these issues: targeting bottlenecks first, automating unplanned or undesirable tasks, tackling high value but deferred work, shifting downstream tasks earlier, and keeping humans in control. Each principle is designed to optimize workflow efficiency and reduce delays, while minimizing pushbacks by positioning AI agents as supportive tools. Our framework enables smoother workflows, faster feedback loops, and improved decision-making, fostering the development of AI agents that deliver measurable productivity gains and drive superior outcomes in real-world applications.

Keywords: *AI agents, productivity boost, workflow optimization, shift-left AI, Lead time reduction, theory-of-constraints, human-in-the-loop*

I. INTRODUCTION

In today's fast-paced work environments, organizations face mounting pressure to improve productivity, streamline processes, and deliver high-quality outcomes under tight deadlines.

AI agents are specialized software systems powered by artificial intelligence which offer significant potential to address these challenges by automating repetitive tasks, generating actionable insights, and optimizing team workflows.

AI agents are particularly effective because they combine machine learning, natural language processing, and rule-based logic to execute tasks with a high degree of awareness about the team's specific context. These could be tasks like drafting documentation, analysing logs, or proposing solutions, all while adapting to evolving project needs. This is accomplished by connecting a Generative AI model with the *tools* already used by the team.

This paper proposes a framework based on five principles for rolling out AI agents for product development teams, in ways that maximize *real* end-end productivity while ensuring team buy-in.

II. LITERATURE REVIEW

Recent research highlights AI agents' potential to boost productivity, with studies showing generative AI like ChatGPT cutting writing task time by 40% and improving quality by 18%[1], and impacting 10% of tasks for 80% of U.S. workers[2].

However, effective deployment requires cost-effective, practical AI agents[3] tailored to real-world needs, as emphasized by critiques and practical guides like Lanham's "AI Agents in Action," which offers frameworks for building production-ready agents[4].

III. AI AGENTS - A LEAN PERSPECTIVE

The principles we propose in this paper draw from *the theory of constraints*, which highlights that a system's performance is capped by its slowest component, or bottleneck.

By focusing AI agents on unblocking bottlenecks, automating tedious or unexpected tasks, tackling high-value but deferred activities, shifting downstream work earlier, and ensuring human oversight, organizations can achieve significant efficiency gains.

This approach aligns with lean methodologies, emphasizing flow efficiency, waste elimination, and continuous improvement. It also incorporates behavioral insights to ensure AI agents are perceived as collaborative partners, maintaining trust and psychological safety within teams. By exploring practical applications such as automating bug triage, generating incremental documentation, or drafting architecture decision records, these principles provide an accessible guide for organizations new to AI agent adoption.

IV. THE FIVE PRINCIPLES

The following sections elaborate on each principle, detailing its productivity benefits, reasons for minimal resistance, and practical use cases for AI agent deployment.

Principle #1: Assist the bottlenecks first

Improving anywhere other than the bottleneck doesn't improve flow or overall productivity. Identify the true bottleneck in your team or system (e.g., QA, design, security reviews) and deploy AI agents to support the bottle-necked roles specifically.

When we assist the non-bottleneck roles, this would result in localised over-production of items. This localised over-production does not improve the end-end productivity of the system. In fact, this over-production can result in increased costs because holding the items that do not reach the customer incurs additional costs.

For example, consider a team that has 5 developers and 2 testers. The team requires all user stories to be tested before calling them 'done'. In this situation, assisting developers using AI agents will result in more stories being developed but not tested.

Though the productivity of developers has increased, this has not resulted in the increase in throughput of the entire system.

In addition to this, all the untested user stories lying around in the codebase will make it more difficult for any future user story to be worked on. This causes errors and eventually, time waste. This time waste is the cost of the prior overproduction at the non-bottleneck role.

Why does this give *real* productivity gains?
Because bottlenecks limit overall throughput, improving them raises the performance of the **entire system**, not just one role. Time saved here compounds across the workflow.

Why does this result in *rare* pushbacks?

- People in bottleneck roles often experience high pressure. These agents act as helpers, reducing that burden.
- Since the AI's outputs still require human validation, employees feel supported but not replaced. They remain in control of quality and direction.
- People in the bottleneck roles already know that there is a huge demand for their work. This makes them less likely to feel that they will be replaced.

Outcomes

- System-wide throughput increases instead of local optimizations[4].
- Bottlenecks are unblocked.
- Delays shrink.
- Handoffs speed up.
- Teams start seeing measurable improvements in lead time, delivery predictability, and quality.

Example AI Agent Use Cases

- Drafting functional test cases and augmenting them based on existing code and requirements, waiting for human review and correction
- Drafting bug reports with relevant logs and screenshots after test execution, for human review and refinement
- Security scanning and compliance automation

Principle #2: Assist with unpleasant or unplanned work

People don't mind getting help with work they hate, especially when it protects their personal time. Automate tasks that interrupt flow or occur unexpectedly outside regular working hours. Let AI take the first pass and then let humans choose whether to follow the AI's solution or not.

Why does this give *real* productivity gains?
Unplanned tasks cause unexpected context switches and reduce flow efficiency. Automating them removes hidden time drains that otherwise eat into deep work.

Why does this result in *rare* pushbacks?
This gives employees more predictable working hours.

Outcomes

- Employee retention
- Lower friction due to reduction in unpleasant and adhoc work

Example AI Agent Use Cases

- Drafting a fix for broken builds or CI/CD failures, and presenting it for human review and approval
- Bug triage and auto-assignment based on logs or history

Principle #3: Use AI agents for high-value work that never gets done

High-cognitive-load activities, such as user story splitting or drafting working agreements, are often avoided due to ambiguity, debate, or interpersonal friction. AI agents can generate actionable options which can act as constructive starting points to facilitate decision-making without imposing final solutions.

Why does this give *real* productivity gains?
Team can pick one of the relevant options suggested by AI and use that as the starting point instead of thinking from the ground up. This saves time.

The artifacts created as part of this process drive better flow of work and improve productivity. For example, having smaller user stories lead to better flow.

Why does this result in *rare* pushbacks?
Content suggested by AI serve as constructive starting points, alleviating mental fatigue and fostering collaborative discussion.

Outcomes

- Unlocking of high-value tasks previously sidelined.
- Human attention is freed for strategic work instead of debates

Example AI Agent Use Cases

- User Story splitting (often avoided due to team disagreements or lack of clarity)
- Proposing multiple UI copy variants for a feature
- Drafting team working agreements

Principle #4: Use AI agents to shift-left the downstream work

Tasks such as writing tests or creating user documentation are often deferred until just before the release, leading to an increase in last minute pressure. If the team draws their workflow on a piece of paper, these tasks appear in the right side.

There is value in shifting these right-side items to the left side of the workflow. AI agents can generate these artifacts continuously as the code gets committed (Example: drafting changes to user documentation based on code-commits) or other early-stage activities (Example: creating functional test cases from user stories before writing code).

This shift-left approach ensures a consistent flow of work and reduces last-minute pressure.

Why does this give *real* productivity gains?
When the issues are found towards the end of the release, this results in heavy rework spanning multiple stages (development, code review, testing, etc.). This principle reduces the resulting heavy rework thereby freeing time to do other productive work. This increases productivity.

Why does this result in *rare* pushbacks?
- Employees don't need to go back and reconstruct what happened weeks ago.
- This avoids addressing large volumes of work under tight deadlines.

Outcomes
- Smoother workflows with a steady stream of deliverables, reducing end-of-cycle pressure
- Faster feedback loops, enabling early identification and resolution of issues
- Reduced downstream burden

Example AI Agent Use Cases
- Drafting user documentation incrementally from user stories and code commits, ensuring alignment with feature development and minimizing last-minute documentation surges.
- Generating release notes from PR titles and commit messages
- Summarizing engineering updates for internal communication

Principle #5: Let humans stay in the driver seat

AI agents deliver maximum value when positioned as tools that enhance, rather than replace, human judgment. By generating inputs, options, and summaries, AI supports experts in making informed decisions without undermining their authority.

Why does this give *real* productivity gains?

AI accelerates exploration and preparation without compromising critical thinking. It lets humans focus on judgment and quality while still reducing total effort.

Why does this result in *rare* pushbacks?
- Employees remain in control.
- AI becomes a thought partner, not a threat. This amplifies their insight instead of replacing their role.

Outcomes
- Maintains trust and psychological safety[5]
- Amplified contributions, driving superior outcomes.

Example AI Agent Use Cases
- Drafting architecture decision records (ADRs) for review
- Summarizing customer feedback for product review
- Suggesting refactoring options for developer's consideration

V. METRICS FOR EVALUATION

In order to measure the effectiveness of the application of these five principles, we need to answer the following questions:

1. How do we know if productivity gain is indeed *real*?
2. How do we know if pushbacks are indeed *rare*?

1. How do we know if productivity gain is indeed *real*?

Most common way of measuring productivity gains is the measurement of *Average Lead Time*[6].

Lead Time is the total time it takes for a work item (e.g., a task, user story, or feature) to move from the moment it is requested or enters the workflow (e.g., added to the backlog or "To Do" column) to the moment it is completed (e.g., delivered or marked as "Done").

$$\text{Lead Time}_i = T_{\text{completion}_i} - T_{\text{request}_i}$$

where:
$T_{\text{completion}_i}$ is the timestamp when the i^{th} work item is marked as completed.

T_{request_i} is the timestamp when the i^{th} work item is first requested or enters the workflow.

For a set of (n) work items with Lead Times {Lead Time_1, Lead Time_2, ... , Lead Time_n} the mean Lead Time is calculated as:

$$\text{Mean Lead Time} = \frac{1}{n}\sum_{i=1}^{n} \text{Lead Time}_i$$

If the productivity gain is indeed real, then we would see the reduction in the Mean Lead Time, over time.

2. How do we know if pushbacks are indeed *rare*?

Most common way of measuring the employee enthusiasm for any initiative is the measurement of *Employee Net Promoter Score (NPS)*.

After the employees start using the AI Agents, each of them is asked a simple question – "*How likely are you to recommend this AI Agent to your friends or colleagues, on a scale of 1 to 10?*"

- Any score up to 6 is treated as a *detractor*.
- 7 and 8 are treated as *passive*.
- 9 and 10 are treated as *promoters*.

We aggregate the responses from all the respondents of the survey and come up with NPS using the following formula.

$$\text{NPS} = \left(\frac{\%\ \text{Promoters} - \%\ \text{Detractors}}{\text{Total Respondents}}\right) \times 100$$

This score can range between -100 to +100.

If the pushbacks are indeed rare, then we would see the increase in the NPS, over time.

3. The Real Acid Test for AI Agent Effectiveness

*Any AI Agent is effective only when it demonstrates decrease in Average Lead Time **AND** increase in Employee NPS over long periods of time.*

This approach balances the productivity with acceptance by the team.

VI. CONCLUSION

The integration of AI agents into team workflows, guided by the five principles outlined in this paper, offers a transformative approach to enhancing productivity while maintaining employee trust and engagement. By targeting bottlenecks, automating unpleasant or unplanned tasks, addressing high-value but deferred work, shift-left mindset, and keeping humans in control, AI agents enable organizations to achieve smoother workflows, faster feedback loops, and improved decision-making. These principles ensure that AI agents act as collaborative tools, amplifying human expertise rather than replacing it, thus fostering rare pushbacks and promoting a culture of innovation. As organizations adopt these strategies, they can unlock significant productivity gains, streamline operations, and power teams to focus on strategic, high-impact work, paving the way for sustainable success in an increasingly competitive landscape.

The evidence suggests that strategic deployment can unlock significant gains, streamline operations, and empower teams. Looking ahead, Sarkar (2025) explores how AI agents will redefine employment in the next tech era, emphasizing the need for frameworks like this to ensure AI enhances rather than disrupts human work[7]. This forward-looking perspective is crucial for organizations aiming for sustainable success in a competitive landscape, balancing innovation with employee well-being.

VII. REFERENCES

[1] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," Science, vol. 381, no. 6654, pp. 187–192, Jul. 2023, doi: 10.1126/science.adh2586.

[2] T. Eloundou, S. Manning, P. Mishkin, and D. Rock, "Generative AI at work," Nat. Bur. Econ. Res. Working Paper, no. 31161, Apr. 2023, doi: 10.3386/w31161.

[3] A. Kapoor et al., "AI agents that matter," arXiv preprint arXiv:2407.01502, Jul. 2024. [Online]. Available: https://arxiv.org/abs/2407.01502

[4] M. Lanham, AI Agents in Action. Shelter Island, NY, USA: Manning Publications, 2024.

[5] S. K and A. Sundararajan, "Enhancing workplace productivity and well-being using AI agents," arXiv preprint arXiv:2501.02368, Jan. 2025. [Online]. Available: https://arxiv.org/abs/2501.02368

[6] G. Thimmapuram, K. Kashyap, and V. Ganesan, Metrics for Agile Product Teams: Illustrated Handbook of Do's and Don'ts. Amazon Media EU S.à r.l., 2021.

[7] T. Sarkar, AI Agents and the Future of Work. New York, NY, USA: Routledge, 2025.