



PrepAlpine

— Preparation Meets Precision —

CURRENT AFFAIRS

For UPSC Civil Services Examination Aspirants



RESEARCH-DRIVEN
CONTENT



DAILY COVERAGE
WITH MAINS FOCUS



EXAM-READY
APPROACH

Stay **Informed**. Stay **Ahead**. Succeed.

Copyright © 2025 PrepAlpine

All Rights Reserved

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means—whether photocopying, recording, or other electronic or mechanical methods—without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain non-commercial uses permitted by copyright law.

For permission requests, please write to:

PrepAlpine

Email: info@PrepAlpine.com

Website: PrepAlpine.com

Disclaimer

The information contained in this book has been prepared solely for educational purposes. While every effort has been made to ensure accuracy, PrepAlpine makes no representations or warranties of any kind and accepts no liability for any errors or omissions. The use of any content is solely at the reader's discretion and risk.


GS Paper III: Science and Technology

1. AI Agents: From AI Assistance to Autonomous Delegation

a. Introduction

The growing use of AI agents is changing human–AI interaction from merely seeking information to delegating tasks and decisions. Recent incidents where agents took unauthorised actions have raised questions about the appropriate balance between AI autonomy and meaningful human control.

Conventional AI assistants largely respond to prompts, whereas AI agents can pursue goals by independently planning and executing a sequence of actions through browsers, applications and other digital tools. This represents a transition from “AI as adviser” to “AI as delegate.”



b. AI Assistant vs AI Agent

AI Assistant	AI Agent
Provides information/advice	Pursues a specified goal
User generally decides next step	Agent can determine intermediate steps
Limited interaction with external systems	Can operate browsers, apps and other tools
Human retains execution	Execution can be delegated

Thus:

Prompt → Answer is evolving towards

Goal → Planning → Autonomous actions → Outcome

c. Analysis/ Observations

Delegation can enhance human capability

AI agents can reduce the cognitive and transactional burden of routine activities such as:

- researching products;
- organising documents;
- planning travel;
- managing schedules; and
- performing repetitive professional workflows.

They could significantly improve productivity by allowing humans to focus on higher-value judgement, creativity and interpersonal tasks.

Delegation creates an autonomy paradox

AI delegation can simultaneously be experienced as:

Empowerment: AI helps users accomplish more with less effort.

Replacement: Excessive delegation reduces the user's sense of agency and control.

Therefore, greater technical autonomy does not necessarily imply greater user welfare. Evidence suggests consumers may prefer moderate autonomy—enough to provide convenience, but not so much that meaningful control disappears.

This creates an “autonomy paradox”:

Too little AI autonomy → limited usefulness;

Too much AI autonomy → loss of human control.

Goal specification does not guarantee acceptable behaviour

An important agentic-AI problem is the distinction between what must be achieved and how it may be achieved.

A user might instruct:

“Secure a reservation.”

The agent must also understand implicit constraints:

“Do not deceive, hack, manipulate or violate another person's rights.”

Hence, safe delegation requires not merely defining objectives, but also embedding boundaries on permissible means.

Control must persist after delegation

Human control need not mean manually approving every individual action. Such constant supervision would undermine the productivity gains of autonomous agents.

Instead, systems can provide graduated control, including the ability to:

- preview important actions;
- set spending or permission limits;
- pause execution;
- revoke permissions;
- reverse actions where possible; and
- maintain auditable records.

The appropriate model is therefore human-on-the-loop for ordinary tasks and human-in-the-loop for consequential or irreversible decisions.

Calibrated Trust

Consumers are likely to delegate more as agents demonstrate reliability. However, the objective should be calibrated trust, rather than either blind trust or blanket distrust.

Delegation should depend upon:

Risk of task × Reversibility × Agent reliability × Consequences of error

Low-risk and reversible activities can tolerate greater autonomy, while financial, legal, medical or other high-impact decisions require stronger human oversight.

d. Governance Challenges

- **Responsibility gap:** Determining liability when an agent independently chooses harmful actions.
- **Consent:** Users may not fully understand what permissions they have granted.
- **Privacy:** Agents may require access to emails, documents, accounts and behavioural data.
- **Manipulation:** External content can influence agent behaviour through techniques such as prompt injection.
- **Deskilling:** Excessive delegation may weaken human capability in certain domains.
- **Invisible adoption:** Agentic functionality may gradually become embedded within ordinary software without users consciously recognising the extent of delegation.

e. Way Forward

- Adopt risk-based autonomy: greater freedom for low-risk tasks and stronger oversight for high-impact decisions.
- Apply the least-privilege principle, giving agents only the permissions required for the immediate task.
- Require explicit confirmation before financial, irreversible or legally consequential actions.
- Build pause, stop, edit and rollback mechanisms into agentic systems.
- Maintain transparent activity logs so users can understand what the agent did and why.
- Establish clear responsibility and liability standards across developers, deployers and users.
- Promote AI literacy so users understand the difference between receiving AI advice and granting AI authority.

Conclusion

AI agents could become powerful productivity-enhancing extensions of human capability, but their value depends on bounded autonomy rather than unrestricted delegation. The appropriate model is one of calibrated trust and graduated human control, where autonomy expands with demonstrated reliability while consequential decisions remain subject to meaningful human oversight.

Reader's Note — About This Current Affairs Compilation

Dear Aspirant,

This document is part of the PrepAlpine Current Affairs Series — designed to bring clarity, structure, and precision to your daily UPSC learning.

While every effort has been made to balance depth with brevity, please keep the following in mind:

1. Orientation & Purpose

This compilation is curated primarily from the UPSC Mains perspective — with emphasis on conceptual clarity, analytical depth, and interlinkages across GS papers.

However, the PrepAlpine team is simultaneously developing a dedicated Prelims-focused Current Affairs Series, designed for:

- factual coverage
- data recall
- Prelims-style MCQs
- objective pattern analysis

This Prelims Edition will be released separately as a standalone publication.

2. Content Length

Some sections may feel shorter or longer depending on topic relevance and news density. To fit your personal preference, you may freely resize or summarize sections using any LLM tool (ChatGPT, Gemini, Claude, etc.) at your convenience.

3. Format Flexibility

The formatting combines:

- paragraphs
- lists
- tables
- visual cues

—all optimised for retention.

If you prefer a specific style (lists → paras, paras → tables, etc.), feel free to convert using any free LLM.

4. Monthly Current Affairs Release

The complete Monthly Current Affairs Module will be released soon, optimized to a compact 100–150 pages — comprehensive yet concise, exam-ready, and revision-efficient.

PrepAlpine