AI-Augmented Research Workflows for Social Scientists

Workshop – Symposium on Behavioral AI in Education, Work, and Decision-Making

Instructor	Rubén Fernández-Fuertes
Date	Friday, April 10, 2026
Duration	∼5 hours (with breaks)
Location	ESSCA School of Management, Paris
Prerequisites	None. No prior programming experience required

Overview

This workshop equips researchers with practical skills to use AI tools in their research workflows. Rather than treating AI as a simple chatbot, participants will learn to work with AI *agents*: systems that can read documents, write and execute code, search the literature, and coordinate with each other on complex tasks. No coding is required from participants; you describe what you need in plain English, and the AI handles the technical execution. The session is hands-on throughout, progressing from first interactions to advanced multi-model workflows.

Schedule

Part 1: From Chatbots to Research Agents 10:00–11:30

What AI can do for your research today

- What are large language models (LLMs) and why they matter for research
- Beyond chatbots: AI agents that can read files, write code, and search the web on your behalf
- Live demo: an agent solving a real research task end-to-end
- Hands-on: setting up the environment and your first agent interaction
- Key insight: you describe tasks in plain English, the AI handles the technical execution
- Important caveat: when to trust AI output and when to verify

— Break (11:30–11:45) —

Part 2: AI for Core Research Tasks 11:45–13:00

Practical workflows for the research lifecycle

- **Reading papers with AI**: from PDFs to structured knowledge
  - Extracting text, tables, equations, and figures from any document (OCR)
  - Summarizing, comparing, and synthesizing across papers
  - Building structured literature reviews from document collections
- **Data work**: cleaning, transforming, and analyzing datasets with AI assistance
  - Tell the AI what you need in plain English: "merge these datasets on country-year"
  - The AI generates and runs analysis scripts (Python/R/Stata) for you

- Hands-on exercises with real social science data and papers

— Lunch Break (13:00–14:00) —

Part 3: Working with Multiple AI Models 14:00–15:15

Why one model is not enough

- Different AI models have different strengths: Claude, GPT, Gemini
- Live demo: coordinated agents tackling a research problem
  - One agent searches the literature, another analyzes data, a third cross-checks the results
- How to make AI models check each other's work
- Hands-on: participants try multi-model coordination on a guided exercise

— Break (15:15–15:30) —

Part 4: Putting It All Together 15:30–16:30

From tools to a research system

- Knowledge management with AI: organizing notes, paper summaries, and research ideas
- Building your personal research pipeline: from paper discovery to reading to notes to writing
- **Open project session**: participants apply the tools to their own research question, with instructor support

Part 5: Wrap-up & Discussion 16:30–17:00

- Ethical considerations: bias, reproducibility, academic integrity
- Q&A and open discussion
- Resources and next steps for continued learning

What Participants Will Take Away

1. A working setup of AI research tools on their laptop
2. Practical experience using AI for literature review, data analysis, and writing
3. Understanding of how to coordinate multiple AI models
4. Experience applying these tools to their own research question
5. A curated list of resources for continued learning

Technical Requirements

- Laptop with macOS or Linux preferred; Windows works but with some limitations
- Internet connection
- **Please create these accounts before the workshop**:
  - Anthropic (Claude): `https://console.anthropic.com` — pay-per-use
  - OpenAI: `https://platform.openai.com` — pay-per-use
  - Google AI Studio: `https://aistudio.google.com` — free tier available

- *Optional*: Ollama for running open-source models locally (free, no API key needed, requires 16 GB+ RAM). We will demonstrate Ollama as a zero-cost alternative, though local models are less capable for complex tasks. Cloud-hosted alternatives are available for a fee
- Don't forget your power adapter — your laptop will be working harder than you

About the Instructor

Rubén Fernández-Fuertes is a PhD candidate in Finance at Bocconi University. His research lies at the intersection of financial economics and natural language processing, using large language models to study how textual information shapes expectations, beliefs, and decision-making in financial markets. He has extensive experience building AI-augmented research workflows, including multi-agent systems, automated literature processing pipelines, and knowledge management tools for academic research.