DEMO

First chapter only

Prompt Engineering for Commerce

Writing Prompts That Process Payments Safely

pragma.vision

PHANTOID.COM

Prompt Engineering for Commerce

Trademark Notice

Google, Google Pay, Google Cloud, and Android are trademarks of Google LLC. Stripe is a trademark of Stripe, Inc. Cloudflare and Cloudflare Workers are trademarks of Cloudflare, Inc. Supabase is a trademark of Supabase, Inc. OpenAI and ChatGPT are trademarks of OpenAI, Inc. Claude is a trademark of Anthropic, PBC. W3C is a trademark of the World Wide Web Consortium. Visa is a trademark of Visa International Service Association. OWASP is a trademark of the OWASP Foundation. Midjourney is a trademark of Midjourney, Inc. Canva is a trademark of Canva Pty Ltd. Etsy is a trademark of Etsy, Inc. Amazon is a trademark of Amazon.com, Inc. All other trademarks are the property of their respective owners.

No Affiliation

This book is an independent publication. It is not authorized, sponsored, or endorsed by any of the companies or organizations whose products or services are mentioned herein.

No Professional Advice

The information in this book is provided for educational purposes only. It does not constitute legal, financial, investment, tax, or other professional advice. Readers should consult qualified professionals for guidance specific to their situation.

Code Examples

Code examples in this book are provided for illustration only. They may not be suitable for production use without additional validation, error handling, and security review.

Published by Pragma Vision LLC

First edition, 2026.

Contents

1

Introduction: When Prompts Handle Money

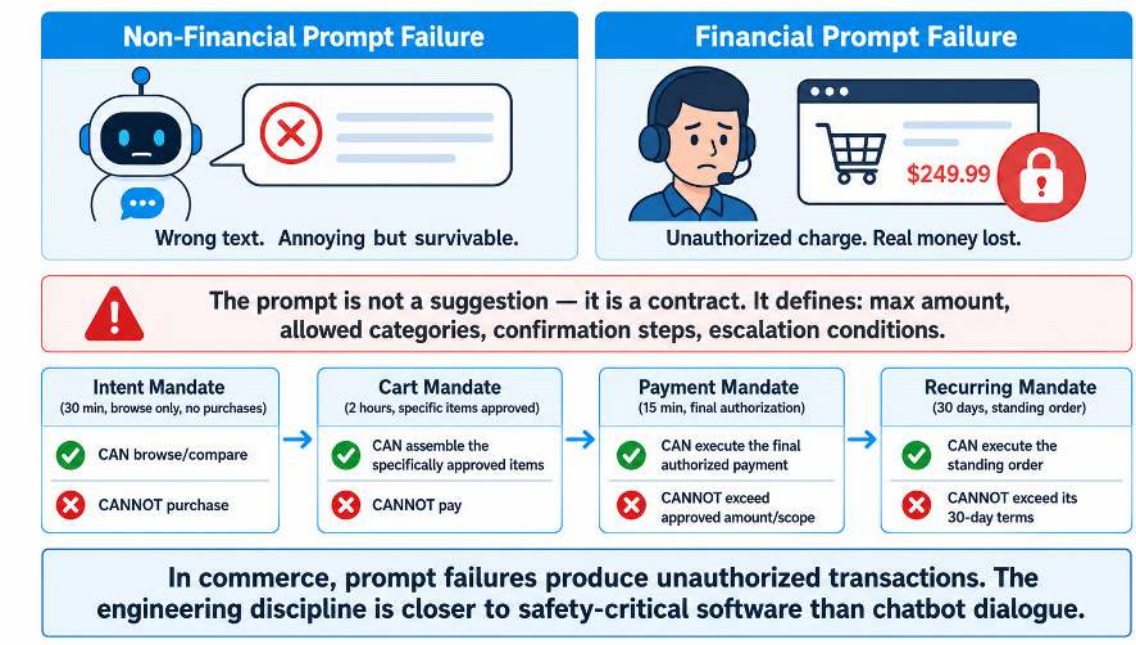

Figure 1. Unauthorized charges are the failure mode that makes commerce prompts contracts, with authority stepping through 30 min intent, 2 hours cart, 15 min payment, and 30 days recurring mandates

1.1 The Prompt Is the Product

Every AI agent is governed by a prompt. The prompt defines what the agent can do, how it reasons, and where it draws the line. For most applications—summarization, coding assistance, creative writing—a mediocre prompt produces mediocre output. Annoying, but survivable. For commerce, a mediocre prompt produces financial losses, unauthorized charges, and broken trust.

The moment an AI agent gains the ability to spend money on behalf of a human, the prompt is no longer a suggestion. It is a contract. It defines the maximum amount the agent can authorize, the categories it can purchase from, the confirmation steps it must follow before committing funds, and the conditions under which it must stop and escalate to the user. A single missing guardrail—one forgotten edge case in the system prompt—can result in an agent authorizing a transaction the user never intended.

73%

of production AI deployments assessed during security audits contained prompt injection vulnerabilities[1]

This book is a practical guide to writing prompts that handle money safely. It is not about generic prompt engineering tips. It is about the specific patterns, guardrails, and testing strategies required when your prompt sits between a user's wallet and the open internet.

Key Insight

In non-financial applications, prompt failures produce wrong answers. In commerce, prompt failures produce unauthorized transactions. The engineering discipline required is closer to writing safety-critical software than writing chatbot dialogue.

[1] OWASP, *Top 10 for Large Language Model Applications*, 2025.

1.2 About Pragma.Vision

Pragma.Vision is an AI-native commerce ecosystem where multiple platforms work together to fulfill human needs through intelligent orchestration. The ecosystem includes a growing family of interconnected platforms—from wish.now (conversational commerce) to phantoid.com (the agent marketplace) to trustauthority.ai (cryptographic identity for AI agents). All transactions flow through three protocol layers: identity verification (Visa TAP), user authorization (Google AP2 with cryptographic mandates), and payment execution (Stripe ACP). Every layer is secured with hybrid signatures combining classical cryptography with quantum-safe ML-DSA-65.

This book draws from the real prompt engineering challenges encountered while building commerce agents for this ecosystem—agents that negotiate prices, process refunds, manage subscriptions, and handle customer escalations across nine platforms simultaneously.

1.3 What You Will Learn

This book covers eight core areas:

1. **Financial Safety Guardrails**: Mandate limits, overspend prevention, confirmation flows, and the architecture of prompts that cannot exceed their authority.
2. **The Purchase Agent**: Product discovery, comparison, and buying prompts with budget awareness and preference learning.
3. **The Negotiation Agent**: Price comparison, bidding strategies, and adversarial resistance for agents that negotiate on your behalf.
4. **The Customer Service Agent**: Refund processing, exchange management, and escalation prompts that balance automation with human judgment.

5. **Prompt Injection Defense**: Attack patterns specific to commerce, prevention techniques, and layered defense architectures.
6. **Testing Commerce Prompts**: Systematic test scenarios, edge cases, adversarial testing frameworks, and continuous monitoring.
7. **Ready-to-Use Templates**: Twenty production-grade prompt templates you can deploy immediately on phantoid.com.
8. **Deployment on phantoid.com**: How to package and deploy your commerce prompts as verified agents on the marketplace.

Whether you are building your first purchase agent or hardening an existing fleet of commerce bots, this book gives you the specific patterns that separate agents that handle money safely from agents that lose it.

DEMO

This is a free preview of the full edition.

Get the complete book at:

https://shop.pragma.vision