



Why Your Payments Data Isn't Ready for AI ... and what you can do about it

Dr. Robert Clewley & Andy Barker
Pagos Solutions

We're sharing a page
from our Pagos
playbook.

How we structure our
approaches to building
reliable data sets for
enterprise business and
get value out of AI.

Going it alone with in-house AI – single PSP

Example application: retry optimization for subscription renewals

You manage one set of:

- decline codes
- scheme fees and descriptors
- dispute attribution and workflow (different reports, formats, etc)
- PSP reports and specs
- timeline reconstruction for transaction lifecycle & dedupe auth rate
- time zone adjustments

AI can probably help you do this...

... But it's not a one-time build

- BIN issuing ranges change weekly
- Product codes update monthly
- Interchange rates change biannually
- PSPs launch new types of service and connection
- Specs can be out of date or ambiguous
- **You might add new PSPs when you grow your business**
 - failover redundancy or routing
 - geographic expansion

Going it alone with in-house AI – multi-PSP

You now must manage:

- Stable customer / order IDs across n PSPs

You can enrich the transactions
in a new data layer

- Comparable view of declines

You now have to manage two sets of
PSP code normalizations and keep
them consistent with each other

Multi-PSP gotchas:

- One reports in ET, another UTC
- One reports amounts in minor units only
- Decline code specifics are different
- One reports only BIN-6
- One PSP drops order ID in metadata
- Another replaces it with their own internal ID
- PSPs might not report useful metadata like issuing bank, country, card type (need accurate BIN lookup)

Part A

Why you're not ready:

1. Fragmented, messy data
2. Missing context

Assumptions:

- you have access to the raw data
- you (will) have multiple PSPs or platforms
- you are busy & under-resourced

Part A

Why you're not ready:

1. Fragmented, messy data
2. Missing context

Part B

How you *can* be ready:

... ?

Generative AI is over-confident, people are skeptical


- Only 8% of businesses are actively using AI for payments today.
- 30% of respondents have **not** identified a payment task they'd be comfortable with AI handling.
- 30% still want Human-in-the-Loop and don't trust it.
- Most teams using AI through PSP-embedded tools, which is not helpful if multi-PSP.

FOMO from hype

**"AI is making the role
of data analyst
obsolete"**

Lived experience

**"I can't trust AI
to take over"**



The Downstream Cost of Bad Data

92% of projects have 1+ small data problems that occur early in a project and compound into large and expensive failures downstream.

45% had 2+.

Reasons?


- "Everyone wants to do the model work, not the data work."
- "Data work is boring, therefore it's easy to automate."


Data Cascades in High-Stakes AI (Google Research, 2021)

<https://dl.acm.org/doi/fullHtml/10.1145/3411764.3445518>


Business data analytics maturity

- ✗ Independent activities
- ✗ Linear progression
- ✗ Further $\nearrow$ is strictly better to be spending time on or aspiring to
- ✗ You can skip to the end if you're behind





The DIKW pyramid




Data enablement is the bottleneck to value



"Everyone wants to do the model work, not the data work."



"Data work is boring, therefore it's easy to automate."



----- Prescriptive
----- Predictive
----- Diagnostic
----- Descriptive
----- Technical



PSP₁

Billing

Fraud/Trust


Networks


PSP₂

CDP

Recovery

Opportunities from multi-PSP data → deeper retry intelligence





PSP₁

Billing

Fraud/Trust

Networks

PSP₂

CDP

Recovery

Strategy

Prescriptive

Revenue and cost optimization

Predictive

Reporting, monitoring, forecasting

Diagnostic

Fraud & dispute handling

Exploratory data analysis

Descriptive

Reconciliation across platforms


Normalization

Technical

- AI can play a **role** as engineer, analyst, strategist

Butterfly effect risks from one-shot AI

- Human has a data warehouse, a business problem, and an LLM subscription
 - Asks LLM to dive in autonomously with a 3-sentence prompt
- Latent issues in data, assumptions, scope
 - Errors magnify in butterfly effect
 - Many divergent outcomes
- Uncertainty about accuracy and relevance
 - Or a complete disaster
 - Harder to audit after the fact



Failure modes

- best case: isolate data for review
- bad case: complete nonsense
- worst case: silently misrepresent data and generate results

Especially tricky with partial or ambiguous source information

- "Bad" specs
- Inherently fuzzy information like BIN mappings and issuer bank names
- Reconciling MID-level vs transaction-level fee reports

1% bad output will be noticed by your customers


Even if AI is accurate 99% of the time,
which 99% is correct?

Failures will be visible frequently:

- many touchpoints in the payments system
- millions of interactions per month

→ 1% frequency is not “rare”

This reduces trust!



Domain expertise

Context layer: business insights, wisdom and strategy

Data Sci. / Analytics

Semantic layer (harmonization) → KPIs
Modeling for classification and prediction

Data Engineering

Data modeling: **ontology** (harmonization)
Operational foundations

Ontology and semantic layer aren't solved problems in CS

“LLMs are remarkably good at producing an initial ontology from documentation and enterprise knowledge.

But every meaningful improvement came from conversations with domain experts. AI can generate structure, **business experts define truth.**”

Need for Human-in-the-Loop guidance


Pagos is constantly updating our mappings:

- AI sometimes starts automatic simple pattern matching
- Validation steps that require careful human review by an expert analyst



Summary – Part A

- AI readiness only possible from data readiness
- AI maturity mirrors DIKW pyramid
- AI can help you work through data readiness
- Expect it to become a full-time job to manage changing data specs
- Use HiL because experts only make it *look* easy



"AI is making the role
of data analyst
obsolete"

"The future of data
analytics isn't fewer
analysts — it's **analysts**
who lead with
judgment, not queries"

Poll

What's your estimate
of your company's
payments data
maturity?


Scale 1-5

1 = very disorganized at all levels;
limited value derived from analytic
opportunities

5 = mature ability to apply business
context and quality data at all levels and
functions; effective at generating value



Part B: Advice for merchants to use LLMs effectively




- **Business questions require a journey back and forth**
- You must help AI act like a human analyst
- You must help AI to be more wise and knowledgeable about your market, industry, biz rules
- You must help AI create "memory" and "learnings" for your role

Where genAI currently shines

- **Working on small, well-scoped problems**
 - Mundane and well-understood tasks
 - Unambiguous data
 - Objective constraints
- **As a companion: "human in the loop"**
 - verify its recommendations vs. allowing it to make decisions
 - maintains an audit trail around decisions
 - documents the steps and concerns
 - provides adversarial review
 - organizes the data being collected
 - accumulates an in-house "memory" of contextual knowledge and rules
 - "semantic layer" + "context layer" + past analysis outcomes

Less uncertainty for Human-in-the-Loop sequence

- Complex analysis problem
 - Human defines and refines at pre-defined, regular milestones
- Latent issues in data, assumptions, scope
 - opportunities to review, learn, realign
- AI owns execution, human owns...
 - intent, judgement, interpretation
- More accuracy, relevance, confidence
 - Audit trail, defensible
 - Iterative improvement for next time



Communications

- Treat like a junior analyst that you directly manage
 - Enthusiastic, well-educated in theory but little practical experience
 - Tone: professional, concise, no emotional or judgemental content
 - skip the “well done”, “please”, and “thank you”
 - avoid criticizing its mistakes
 - Provide onboarding, context, coaching, oversight, feedback, and vetting
 - Expect them to show their work, including citations
 - You are always responsible for the final output

Process checklist for AI collaboration

- Discuss your intent
- Set a goal
- State / discuss assumptions
- **Get feedback and clarifying questions; potential risks / blind spots**
- **Science: observations, hypotheses, tests, review, iterate**
- Maintain README that summarizes the planned approach
- Require output saved in CSVs and markdown reports
- Update templates & system prompts for next time

Summary – Part B

Communication:

- Expectation = smart junior analyst that you manage
- Model what you want to see
- Propose vs assert

Process:

- Collaborate, don't one-shot
- Adversarial review
- Use the scientific method
- Iterate on your setup

Organization:

- Work in a shared repository (keep data separate)
- One project per folder

Low-tech "orchestration":

- Templates, system prompt
- Repo file "memory"
- Glossary of terms
- MCP data access
- No need for agents (yet)

Example snippet of claude.md

When presenting the results of analysis, I want clarity on what are direct observations versus hypothetical explanations. Do not present "likely reasons" as fact. Use caution and skepticism like an analyst or scientist should, and consider multiple ways in which results might be wrong, invalid, or biased.

Consider running additional checks and validations in the spirit of seeking falsification. Use reasonable metrics for determining confidence or consult with the user before going on long tangents.

For new areas of investigation, let me review your strategy to query the data before you run it.

Further reading

- <https://ebizcharge.com/blog/ai-in-payment-processing/>
- <https://career.uml.edu/blog/2025/08/29/will-data-analysts-be-replaced-by-ai/>
- <https://www.databricks.com/blog/future-data-analytics-why-ai-rewriting-analysts-job-description>
- https://www.linkedin.com/posts/msmullins_my-feed-is-full-of-ontology-chatter-and-im-activity-7477004838145396737--tsv
- <https://mollyg.substack.com/p/lazy-fucking-interns>
- <https://lnkd.in/p/eQsfdUDQ>
- <https://www.tellius.com/resources/blog/why-your-semantic-layer-isnt-ready-for-ai-agents-and-what-to-do-about-it> [5-level AI analytics maturity]
- <https://atlan.com/know/context-layer-vs-semantic-layer/>
- https://medium.com/@grom_65116/the-semantic-layer-is-dead-now-its-an-api-for-ai-agents-f91d48a0c74a
- <https://www.seangoedecke.com/llms-reward-expertise>