Specimen Exam 1

Exam Name: Data Science Professional (SDS)
Time Allowed: 3 hours 15 minutes
Pass Mark: 50%

This examination contains 2 sections:

Section A:
1 question, worth 50 marks

Section B:
2 questions, each worth 25 marks

Total 100 marks

Note: This examination contains a synthetic dataset, its creation was assisted by Generative AI and reviewed by ACCA.

Section A

Question 1

The following **exhibits** provide information relevant to the question:

1. Organisation information – the background, industry information and current operations for the company: EssanPay Co (EP)
2. Predictive model – details of a predictive model developed to predict fraudulent transactions
3. Model underperformance – details of how the model has underperformed relative to expectations
4. Risk concerns – details of recent difficulties faced by a competitor, prompting the CRO to assess how EP can avoid similar risks
5. Appendix 1 – file describing the features used to train the predictive model
6. Appendix 2 – file containing detailed information on EP's payment transactions during January 2026

1 Organisation information

EssanPay Co (EP) is a mid-sized payments and digital financial services company operating in Zeeland. EP provides card payment processing and digital wallets to retail and online merchants across several markets.

A key part of EP's operations is the real-time identification of potentially fraudulent transactions.

In recent months, EP has experienced rapid growth in online and mobile transactions, alongside increased adoption of new customer authentication technologies. EP operates in a highly regulated environment, and senior management places strong emphasis on operational resilience, customer experience, and effective risk management.

2 Predictive model

The data science function developed a Naïve Bayes classification model to identify potentially fraudulent transactions. The model has been in operation for several months now.

The model predicts outcomes for each transaction as follows:

- Legitimate = 0
- Fraudulent = 1

Transactions predicted as fraudulent are routed to a specialist fraud operations team for manual review. Transactions predicted as legitimate are processed automatically.

The classification model was trained using the confirmed fraud outcomes in respect of approximately 5,000 historical transactions over a six-month period. An accuracy rate of 97.8% was recorded prior to the model's deployment.

Appendix 1 contains a data dictionary which describes the features used to the train the model.

Following a review of the most recent quarter, the chief risk officer (CRO) has raised serious concerns that:

- A significant number of fraudulent transactions were not identified by the model, resulting in direct financial losses, increased chargebacks, regulatory risk and reputational risk.
- A large number of legitimate transactions were incorrectly flagged as suspicious, leading to customer annoyance, payment friction, and increased workload for the fraud operations team.

Details of the transactions implemented in January 2026 are included in Appendix 2. Appendix 2 contains the following relevant fields:

- actual_fraud (0 = Legitimate, 1 = Fraudulent)
- predicted_fraud (0 = Legitimate, 1 = Fraudulent)

The following confusion matrix has been produced by the data science function to evaluate the model's performance:

	Predicted Legitimate	**Predicted Fraudulent**
Actual Legitimate	1,216	131
Actual Fraudulent	23	7

3 Model underperformance

The CRO is particularly disappointed in the outcome, as the predictive model had demonstrated a very high accuracy level in training. He has asked the head of the data science function how this could have occurred.

Following an investigation, the head of the data science function has suggested four reasons for the model underperformance. These are class imbalance, data quality issues, overfitting, and model drift.

4 Risk concerns

The CRO has recently seen press coverage concerning a competitor, FernLink Co.

The share price of FernLink Co has experienced a very sharp decline. The share price decline has been attributed to significant financial losses and reputational damage. The financial losses relate to an increase in the volume of fraudulent transactions incurred. FernLink Co has also suffered reputational damage following a regulatory investigation into its fraudulent transactions and a spike in customer complaints due to the number of legitimate transactions flagged as suspicious. It is suggested that governance of their fraud prediction models has contributed significantly to this.

The CRO wishes to understand how EP can mitigate against similar issues.

5 Appendix 1

	A	B	C
1	Column	Type	Description
4	customer_id	String	Synthetic customer identifier.
5	customer_age	Integer	Customer age in years
6	customer_tenure_days	Integer	Days since customer signup
7	is_vip	Binary (0/1)	1 if customer is VIP / high-value segment.
8	merchant_id	String	Synthetic merchant identifier.
9	merchant_category	Categorical	Merchant sector/category (e.g., grocery, electronics, travel, etc.).
10	merchant_risk_score	Float (0–1)	Merchant risk propensity (higher = riskier).
11	amount	Float	Transaction amount (positive; lognormal-like).
12	currency	Categorical	Transaction currency (EUR/GBP/USD).
13	payment_method	Categorical	Payment rail (card/bank_transfer/wallet).
14	channel	Categorical	Origin channel (web/mobile_app/pos; may include in_app in production).
15	card_present	Binary (0/1)	1 if card-present (mostly POS).
16	device_type	Categorical	Device OS (ios/android/windows/macos/linux).
17	device_fingerprint	String / Null	Pseudo device identifier; may be missing in production.
18	billing_country	Categorical	Billing country code.
19	shipping_country	Categorical	Shipping country code (may differ from billing).
20	ip_country	Categorical	IP geolocation country code
21	is_cross_border	Binary (0/1)	1 if billing/shipping/IP countries are inconsistent (proxy for cross-border risk).
22	txn_count_1h	Integer	Number of transactions for customer in last 1 hour (velocity feature).
23	txn_count_24h	Integer	Number of transactions for customer in last 24 hours (velocity feature).
24	prev_chargebacks_12m	Integer	Number of prior chargebacks in last 12 months.
25	email_domain	Categorical	Customer email domain (e.g., gmail.com, business.com, student.edu, ...).
26	is_new_device	Binary (0/1)	1 if device appears new for this customer (proxy).
27	actual_fraud	Binary (0/1)	Target label: 1 if fraudulent, 0 otherwise.

6 Appendix 2

Refer to the spreadsheet titled SDS Data Science Professional Specimen Examination 1 – Question 1 Appendix 2

Requirements:

You are an accountant in EssanPay Co who has been seconded to the data science function.

Write a report for the chief risk officer (CRO) of EssanPay Co which:

(i) Evaluates the performance of the predictive model using the transaction data for January 2026 and the confusion matrix provided;

(12 marks)

(ii) Evaluates the four reasons suggested by the head of the data science function for the model's underperformance;

(18 marks)

(iii) Recommends the actions which could be taken by the data science function to improve the governance of future fraud prediction models.

(10 marks)

Professional marks will be awarded for the demonstration of skill in communication, analysis and evaluation, scepticism and commercial acumen in your answer.

(10 marks)

(50 marks)

Section B

Question 2

The following **exhibits** provide information relevant to the question:

1. Organisation information – the background, industry and structure of the company: Watlee Co
2. SQL code - SQL table definitions proposed by a junior analyst for a reporting area
3. Dashboard discussion – details of a management conversation about the dashboard used for reporting purposes
4. Appendix 1 – the dashboard produced by the costing department for the chief financial officer

1 Organisation information

Watlee Co (Watlee) is a mid-sized manufacturing company producing packaged dessert products for retail and hospitality customers across multiple regions. The company manufactures six products across four manufacturing plants, each with different cost structures and production efficiencies.

Products incur direct material and labour costs, while overhead costs (such as utilities and plant administration) are incurred at plant level and allocated to products using machine hours. Production volumes, machine usage, and costs vary by product, plant, and month.

Watlee produces management reports on a monthly basis.

2 SQL code

Operational and financial data is stored across multiple systems at Watlee, including production, costing, machine usage, overhead, and sales systems. These datasets are periodically joined and aggregated to support management reporting.

Below are some draft reporting tables produced by a junior analyst using SQL data types:

```
-- Draft reporting tables (junior analyst)

-------------------------------------------------
-- Production table
-------------------------------------------------

CREATE TABLE stg_production (
  plant_id        VARCHAR(10) NOT NULL,
  product_sku     TEXT NOT NULL,
  prod_month      VARCHAR(7) NOT NULL,  -- store month e.g., '2026-01'
  units_produced  FLOAT,
  machine_hours   FLOAT,

  -- Primary Key
  CONSTRAINT pk_stg_production
    PRIMARY KEY (plant_id, product_sku, prod_month)
);

-------------------------------------------------
-- Costs table
-------------------------------------------------

CREATE TABLE stg_costs (
  plant_id            VARCHAR(10) NOT NULL,
  product_sku         TEXT NOT NULL,
  prod_month          VARCHAR(7) NOT NULL,
  material_unit_cost  FLOAT,
  labour_unit_cost    FLOAT,

  -- Primary Key
  CONSTRAINT pk_stg_costs
    PRIMARY KEY (plant_id, product_sku, prod_month),

  -- Foreign Key to production
  CONSTRAINT fk_costs_production
    FOREIGN KEY (plant_id, product_sku, prod_month)
    REFERENCES stg_production (plant_id, product_sku, prod_month)
);

-------------------------------------------------
-- Overhead table
-------------------------------------------------

CREATE TABLE stg_overhead (
  plant_id        VARCHAR(10) NOT NULL,
  prod_month      VARCHAR(7) NOT NULL,
  overhead_cost   FLOAT,

  -- Primary Key
  CONSTRAINT pk_stg_overhead
    PRIMARY KEY (plant_id, prod_month)

);
```

3 Dashboard discussion

Watlee's chief financial officer (CFO) is considering how to improve the presentation of management information to support monthly performance discussions and has recently met with various stakeholders across the business to review a dashboard which has been produced by the costing department.

Below is a transcript of the discussion:

CFO:

"We want this dashboard to become the primary tool used at executive meetings. It should quickly show whether financial performance is improving or deteriorating and allow management to identify areas requiring action.

Head of costing:

"The layout looks clear, but I'm not sure it gives us the level of detail we need. From a costing perspective, we normally analyse performance by cost type, for example materials, labour and overhead, rather than just total cost of sales. Also, if costs are increasing, we need to understand why. For example, whether this is driven by input prices, efficiency issues, or production mix changes."

Commercial manager:

"I'm also wondering whether the dashboard shows a complete financial picture. I think it's difficult to judge overall performance from the dashboard. The dashboard should present useful and complete financial information, not just activity measures."

Data analyst:

"We should also consider whether this meets accessibility requirements, particularly for users with colour-vision deficiencies."

Operations manager:

"One concern I have is that charts can sometimes be misleading depending on how data is presented. We need confidence that managers won't draw incorrect conclusions. Also, there is no supporting commentary explaining what the dashboard is telling us. Senior management may interpret results differently unless there is some narrative or key insights alongside the visuals."

CFO:

"So to summarise, when we review this dashboard at executive meetings, we need to be confident that it presents reliable information, clearly highlights areas requiring action and supports consistent interpretation across management. I am going to organise an independent review of how effectively this dashboard supports executive decision-making before we adopt it fully."

4 Appendix 1

Executive Summary

Revenue

Month	Actual	Budget	Var %	Diff
January	$1,098K	$1,228K	-11%	(130K)
February	$1,079K	$1,329K	-19%	(250K)
March	$1,295K	$1,385K	-6%	(90K)
April	$1,279K	$1,085K	18%	194K
May	$1,279K	$1,085K	18%	194K
June	$1,279K	$1,545K	-17%	(266K)
July	$1,313K	$1,505K	-13%	(192K)
August	$1,329K	$963K	38%	366K
September	$1,399K	$1,655K	-15%	(256K)
October	$1,399K	$1,789K	-22%	(390K)
November	$1,295K	$1,799K	-28%	(504K)
December	$1,639K	$1,719K	-5%	(80K)
Total	**$15,683K**	**$17,087K**	**-8%**	**(1,404K)**

Cost Of Sales

Month	Actual	Budget	Var %	Diff
January	$267K	$357K	-25%	(90K)
February	$334K	$277K	21%	57K
March	$315K	$224K	41%	91K
April	$316K	$279K	13%	37K
May	$219K	$223K	-2%	(4K)
June	$261K	$268K	-3%	(7K)
July	$299K	$269K	11%	30K
August	$236K	$239K	-1%	(3K)
September	$298K	$399K	-25%	(101K)
October	$399K	$399K	0%	0K
November	$319K	$385K	-17%	(66K)
December	$351K	$331K	6%	20K
Total	**$3,614K**	**$3,650K**	**-1%**	**(36K)**

Unit Cost by Plant 2025 2026

Unit Cost by Product

Product W	$4.80
Product X	$4.50
Product Y	$4.00
Product Z	$3.60

Revenue Actual Versus BudMonth

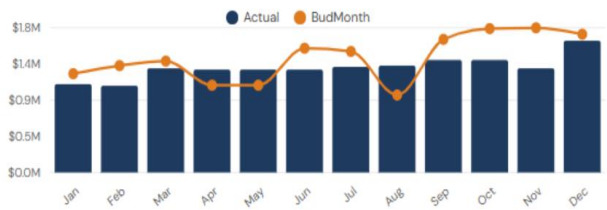

Cost Of Sales Actual Versus Budget

Revenue Trend by Plant

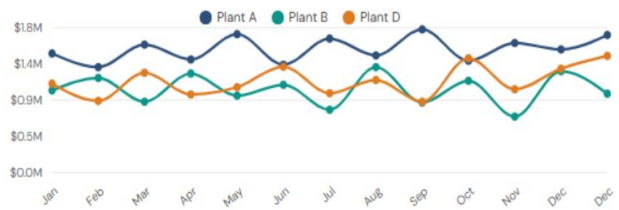

Requirements:

(a) Evaluate the suitability of the SQL data types used in the proposed tables for Watlee Co's reporting requirements and recommend improvements to the data types and table design.

(10 marks)

(b) Evaluate how effectively the dashboard supports executive decision-making, taking account of the stakeholder discussion.

(10 marks)

Professional marks will be awarded for the demonstration of skill in analysis and evaluation, and commercial acumen in your answer.

(5 marks)

(25 marks)

Question 3

The following **exhibits** provide information relevant to the question:

1. Organisation information – the background, industry and operations of the company: MoyWorth Co (MW)
2. CRISP-ML project plan – information related to a proposed data science project
3. Stakeholder management – details of key stakeholders identified for the proposed data science project
4. Appendix 1 – excerpt from the CRISP-ML project plan submitted by the junior data scientist

1 Organisation information

MoyWorth Co (MW) is a regulated financial services organisation seeking to enhance its financial reporting processes using data science techniques. The organisation prepares quarterly management accounts and annual statutory financial statements in line with applicable accounting standards and regulatory requirements, operating under strict internal and external reporting deadlines.

Financial data are sourced from multiple systems, including policy administration platforms, claims systems, investment platforms, and the central finance ledger. These systems have been developed or acquired at different times and use differing data definitions and reporting conventions.

As a result, the reporting process involves significant manual adjustment and reconciliation, often carried out under time pressure close to reporting deadlines. Senior management is concerned that adverse trends in profitability and capital adequacy are frequently identified too late in the reporting cycle, limiting the scope for effective management action.

Management has expressed a desire to incorporate more forward-looking insight into the reporting cycle. The organisation holds several years of detailed historical financial and operational data. Preliminary internal analysis suggests that adverse outcomes are influenced by complex interactions across multiple datasets. These interactions are difficult to detect using traditional rule-based analysis.

2. CRISP-ML project plan

To address this issue, a data science project using machine learning techniques was initiated. The first three stages of the project plan were completed by a senior data scientist to a very high professional standard using the CRISP-ML framework. Following the unexpected departure of the senior data scientist, a junior data scientist was tasked with completing the remaining stages of the project plan using CRISP-ML.

The plan submitted by the junior data scientist to the chief financial officer (CFO) is in Appendix 1.

3. Stakeholder management

The proposed data science project will involve multiple business and technical stakeholders. The model is intended to provide forward-looking insight to help management anticipate adverse trends and take action earlier in the reporting cycle.

Internal stakeholders have expressed varying expectations regarding the balance between innovation, governance oversight and operational reliability.

The following key internal stakeholders have been identified:

Finance leadership team (FLT)

The FLT is accountable for the accuracy of financial reporting and for meeting internal and regulatory reporting deadlines. The team expected to use insights generated by the model to support a timelier investigation of emerging risks and performance issues.

The FLT is primarily concerned with whether model outputs are understandable and actionable, avoiding disruption to established reporting processes and ensuring insights are delivered in a timely and practical manner.

Some finance managers have indicated caution regarding reliance on analytical models which they do not fully understand.

Risk management department (RMD)

The RMD is responsible for governance oversight, internal controls, and ensuring that new analytical tools are used appropriately within the organisation’s risk framework.

Its focus includes ensuring appropriate governance and documentations, preventing inappropriate reliance on model outputs and maintaining compliance with regulatory expectations relating to risk management and decision-making processes.

The RMD has emphasised the importance of transparency and clear accountability before operational use of the model.

IT & data engineering teams

IT and data engineering teams are responsible for system stability, data pipelines and ongoing operational maintenance of analytical solutions.

Their priorities include ensuring solutions integrate with existing infrastructure, maintaining data quality and system reliability and managing resource constraints and operational support responsibilities.

IT representatives have noted that previous analytical initiatives experienced challenges when operational ownership and maintenance responsibilities were unclear.

4. Appendix 1

Stage 4 – Model Evaluation

The model will be evaluated to ensure it performs adequately before being used by the Finance teams.

Planned actions include:

- Calculating accuracy metrics
- Comparing predicted results with historic financial results
- Adjusting the model where performance appears unsatisfactory.

If the model results appear reasonable and accuracy is acceptable, the model will proceed to deployment.

Stage 5 – Model Deployment

Deployment activities include:

- Publishing the dashboards to users
- Scheduling automatic data refreshes
- Running the model prior to reporting deadlines to generate updated outputs.

Finance teams will then incorporate insights from the dashboards into reporting processes.

Stage 6 – Monitoring and Maintenance

After deployment, the model will be reviewed periodically to ensure it continues to operate as expected.

Planned actions include:

- Reviewing results periodically
- Updating the model when required
- Gathering informal feedback from Finance users and adjusting the model if necessary.

Requirements:

(a) Evaluate the junior data scientist's project plan in respect of the CRISP-ML stages for which they were responsible (Stages 4-6) and recommend improvements to the project plan to address any weaknesses identified.

(12 marks)

(b) Advise on how the stakeholders identified should be managed throughout the project lifecycle.

(8 marks)

Professional marks will be awarded for the demonstration of skill in analysis and evaluation, scepticism, and commercial acumen in your answer.

(5 marks)

(25 marks)