

Quiz: Lesson 7.4 – Model Risk Governance and Validation
Module 7: The Compliance Problem

Prof. Dr. Joerg Osterrieder

Question 1 (Understand)

What is the definition of model risk according to SR 11-7?

- A The risk that a model will be too expensive to maintain
- B **The potential for adverse consequences from decisions based on incorrect or misused model outputs**
- C The probability that a model will produce statistically insignificant results
- D The risk that competitors develop better models

Question 2 (Understand)

Which regulatory document established the foundational framework for model risk management in US banks?

- A Basel III Capital Accord
- B GDPR Article 22
- C **SR 11-7 (2011)**
- D Dodd-Frank Act Section 165

Question 3 (Understand)

What are the three pillars of the SR 11-7 framework?

- A Risk assessment, capital allocation, reporting
- B **Model development, model validation, model use and governance**
- C Data quality, algorithm design, deployment
- D Prevention, detection, response

Question 4 (Understand)

In the three lines of defense, what is the role of the 2nd line?

- A Build and maintain models in production
- B **Provide independent validation and effective challenge of models**
- C Audit the overall model risk governance framework
- D Deploy models into production systems

Question 5 (Understand)

Which of the following is NOT one of the three components of model validation?

- A Conceptual soundness
- B Outcome analysis
- C Benchmarking
- D **Profitability assessment**

Question 6 (Apply)

A large bank has a Basel IRB PD model used for regulatory capital calculations. What tier should this model be assigned?

- Tier 1 (annual validation, full independent review)**
- Tier 2 (biennial validation, targeted review)
- Tier 3 (triennial validation, streamlined review)
- Untiered (no formal validation required)

Question 7 (Apply)

A bank's VaR model produced 7 exceedances over 250 trading days at the 99% confidence level. In which Basel traffic light zone does this fall?

- A Green zone (no capital impact)
- B **Yellow zone (capital multiplier of 1.38x)**
- C Red zone (capital multiplier of 1.75x or higher)
- D Cannot be determined without additional information

Question 8 (Understand)

What is the primary purpose of a model inventory?

- A To track the profitability of each model
- B To store the source code of all models
- C **To maintain a centralized register of every model used, including ownership, validation status, and tier**
- D To rank models by computational performance

Question 9 (Analyze)

In the London Whale incident (2012), which model governance failure was most directly responsible for the \$6.2 billion loss?

- A The model used incorrect market data feeds
- B The model was too complex for the available compute infrastructure
- C **The VaR model was changed without independent validation or formal change control**
- D The model had never been backtested

Question 10 (Apply)

During champion-challenger testing, the challenger model shows a Gini coefficient of 0.72 compared to the champion's 0.70, but the challenger uses a deep neural network while the champion is logistic regression. What should the Model Risk Committee recommend?

- A Immediately replace the champion with the challenger
- B **Extend the test period, as the marginal improvement may not justify the added complexity and explainability challenges**
- C Discard both models and develop a new one
- D Deploy both models simultaneously in production

Question 11 (Understand)

What does the Population Stability Index (PSI) measure in model monitoring?

- A The accuracy of model predictions over time
- B **Whether the distribution of model inputs or outputs has shifted compared to the development sample**
- C The computational stability of the model under load
- D The percentage of the population covered by the model

Question 12 (Analyze)

A model validator finds that a credit risk model's Gini coefficient has dropped from 0.65 to 0.52 over the past year. What type of drift does this indicate and what action is required?

- A Data drift; no action needed until the next scheduled review
- B **Performance drift exceeding the 10% threshold; escalation to Model Risk Committee is required**
- C Concept drift; the model should be immediately retired
- D Parameter drift; recalibrate the model and continue

Question 13 (Understand)

Why is non-determinism a fundamental challenge for GenAI/LLM governance?

- A It makes models too slow for real-time applications
- B It increases the computational cost of model execution
- C **It means the same input can produce different outputs, breaking traditional validation approaches that assume reproducibility**
- D It prevents models from being deployed in cloud environments

Question 14 (Apply)

A bank wants to use an LLM to generate summaries of credit committee meeting notes. Under current best practices, how should this be classified?

- A Tier 3 (low risk, minimal governance needed)
- B Not a model (no governance required)
- C **Tier 1 (high risk by default for GenAI, requiring full governance including human-in-the-loop review)**
- D Tier 2 (medium risk, standard governance)

Question 15 (Understand)

What is prompt injection in the context of LLM governance?

- A A technique for improving model accuracy by providing better prompts
- B **An attack where adversarial instructions are embedded in data to manipulate LLM behavior**
- C A method for calibrating LLM temperature parameters
- D A standard validation technique for testing LLM robustness

Question 16 (Apply)

An audit trail for a model must be retained for a minimum of how many years, and in what type of storage?

- A 3 years, standard database
- B 5 years, encrypted cloud storage
- C **7 years (10 in some jurisdictions), immutable write-once-read-many (WORM) storage**
- D 1 year, any storage with backup

Question 17 (Analyze)

A sensitivity analysis reveals that a 0.5% change in a single input parameter causes a 40% change in model output. What does this indicate?

- A The model is well-calibrated and responsive to market conditions
- B The model is performing within expected parameters
- C **The model has a cliff effect and may be fragile — this should be flagged as a validation finding**
- D The model needs more training data to stabilize

Question 18 (Understand)

Which UK regulation, effective from 2024, established comprehensive model risk management principles for PRA-regulated firms?

- A FCA Handbook SYSC 7
- B Bank of England Monetary Policy Framework
- C **PRA SS1/23**
- D UK Online Safety Act

Question 19 (Evaluate)

A bank is deciding whether to use a complex ensemble ML model or a simpler logistic regression for credit scoring. The ML model has a Gini of 0.74 and the logistic regression has a Gini of 0.71. Which is the better choice from a model risk governance perspective?

- A The ML model, because it has higher predictive power
- B Either model is acceptable since both exceed a 0.70 Gini threshold
- C **The logistic regression is likely preferable, as the marginal accuracy gain does not justify the added complexity, explainability challenges, and governance costs**
- D Neither model is acceptable; a Gini of at least 0.80 is required for regulatory models

Question 20 (Evaluate)

Under GDPR Article 22 and the EU AI Act, what specific audit trail capability is required for AI-driven credit decisions?

- A Logging the final credit score only
- B Recording the model name and version number
- C **Providing a meaningful explanation of the decision logic, including the key factors that influenced the outcome, retrievable for the affected individual**
- D Storing a copy of the entire training dataset