

Contents

| | |
|---|----------|
| Chapter 7 — Intuitive MCQ Bank (Legal and Ethical) | 1 |
| Answer Key | 16 |
| Answer Distribution | 16 |

Chapter 7 — Intuitive MCQ Bank (Legal and Ethical)

56 multiple-choice questions, intuitive/conceptual level. One correct option each. Cover the answer (blockquote) while testing yourself. Bilingual explanations (English + বাংলা).

Exam style: “Which statement best describes...”, exactly one correct option; use full technical terms, no abbreviations.

বাংলা ব্যাখ্যা: এই ব্যাংকে অধ্যায় ৭-এর ৫০+ ধারণামূলক MCQ আছে প্রতিটি প্রশ্নের নিচে উত্তর ও সংক্ষিপ্ত ব্যাখ্যা (ইংরেজি + বাংলা)

Topic: Foundations — Why Responsibility Matters

Q1. The lecture states that “a system can work exactly as intended — and still cause harm.” Which scenario best illustrates this idea?

- A. An engagement-maximizing recommender works perfectly and amplifies outrage content, radicalizing some users.
- B. A model crashes with a null-pointer exception during inference and returns no output.
- C. A developer accidentally swaps two labels in the training set and the accuracy drops.
- D. A server outage prevents users from reaching the deployed model for an hour.

Answer: A. The harm comes from the system doing precisely its job (maximizing engagement), not from a malfunction; B, C, and D are ordinary technical failures or bugs, which is the opposite of “working as intended.” **বাংলা:** সিস্টেম নিখুঁতভাবে কাজ করেও ক্ষতি করছে — এটাই মূল বার্তা; বাকিগুলো নিছক বাগ বা ত্রুটি

Q2. Why does the chapter argue that errors in AI systems become “systemic” rather than isolated?

- A. Because AI code is always written in interpreted languages that hide bugs.
- B. Because a single biased weight pattern repeats the same mistake across millions of decisions.
- C. Because AI systems are legally exempt from product-liability rules.
- D. Because every AI system is deployed by exactly one party who controls everything.

Answer: B. One model makes vast numbers of decisions, so a flawed pattern is reproduced everywhere at once; A is irrelevant, C is false (liability now extends to software/AI), and D contradicts the “distributed responsibility” point. **বাংলা:** একটাই মডেল লক্ষ লক্ষ সিদ্ধান্ত নেয়, তাই একটা পক্ষপাত সর্বত্র একসাথে ছড়ায়

Q3. The chapter identifies three legal stakeholders of every AI system. Which trio is correct?

- A. The programmer, the tester, and the project manager.
- B. The shareholder, the regulator, and the auditor.
- C. The data subject, the user/affected person, and the third party (copyright/likeness holder).
- D. The cloud provider, the data center, and the network operator.

Answer: C. The chapter names the data subject (whose data is used), the affected person (judged by outputs), and the third party (whose copyright or likeness may be infringed); the others are internal or infrastructural roles, not the legal stakeholder triad. **বাংলা:** তিন স্টেকহোল্ডার — যার ডেটা, যে সিদ্ধান্তের শিকার, আর যার সৃষ্টিকর্ম/অবয়ব ব্যবহৃত হয়

Q4. Why does the lecture call data choices “legal and ethical design decisions”?

- A. Because data is stored in legally protected file formats.
- B. Because only lawyers are permitted to choose training data.
- C. Because design decisions are reversible at any time after deployment.

- D. Because each design decision (what to collect, what objective to optimize, what fallback to provide) changes the exposure of each stakeholder.

Answer: D. The chapter ties every engineering choice to the risk borne by data subjects, affected persons, and third parties; A is irrelevant, B is false, and C contradicts the difficulty of undoing past training. **বাংলা:** প্রতিটি ডিজাইন-সিদ্ধান্ত কোনো না কোনো স্টেকহোল্ডারের ঝুঁকি বাড়ায়/কমায়, তাই এগুলো আইনি-নৈতিক সিদ্ধান্ত

Q5. Which set best captures the “engineering consequences” the chapter says follow from distributed responsibility?

- A. Clear ownership of each pipeline stage, traceability, and monitoring with a safe fallback.
- B. Faster hardware, lower latency, and cheaper storage.
- C. Larger models, more parameters, and longer training runs.
- D. Marketing transparency, brand consistency, and customer loyalty.

Answer: A. Because responsibility is split across data, training, integration, and operation, the chapter demands accountable ownership, traceability (logs, versioning), and monitoring/fallback; the other options are performance or business goals, not responsibility safeguards. **বাংলা:** দায় ভাগ হয়ে যায় বলে দরকার স্পষ্ট মালিকানা, ট্রেসেবিলিটি আর মনিটরিং+ফলব্যাক

Topic: Why AI Is Not “Just Software”

Q6. Which statement best describes the core difference between traditional software and an AI system?

- A. Traditional software is probabilistic, while AI systems are deterministic.
- B. In traditional software behavior is fully specified by the developer, whereas in an AI system decisions are learned, not fully specified.
- C. AI systems never change behavior over time, while traditional software drifts.
- D. Traditional software ships behavior, while AI systems only ship code.

Answer: B. The defining contrast is specified versus learned behavior; A reverses determinism, C reverses drift (AI drifts), and D inverts the lecture’s slogan “you ship behavior, not code.” **বাংলা:** সাধারণ সফটওয়্যারে আচরণ স্পষ্ট-নির্দিষ্ট, AI-তে আচরণ শেখা হয় — এটাই মূল পার্থক্য

Q7. A unit test `assertEqual(answer, "Funafuti")` is described as nearly useless for a language model. Why?

- A. Because string comparison is computationally too expensive at scale.
- B. Because language models cannot output proper nouns.
- C. Because many phrasings can be valid and the model is probabilistic and may occasionally hallucinate, so one exact-match check cannot capture correctness.
- D. Because unit tests are forbidden by the EU AI Act.

Answer: C. The output is open-ended and non-deterministic, so a single exact-string assertion fails to test real behavior; A is false (comparison is cheap), B is absurd, and D invents a rule. **বাংলা:** বহু সঠিক উত্তর সম্ভব এবং মডেল probabilistic, তাই একটামাত্র exact-match পরীক্ষা অর্থহীন

Q8. Given that AI behavior is learned and probabilistic, which quality-assurance approach does the chapter recommend?

- A. A single deterministic regression test that must pass exactly every time.
- B. Replacing all automated tests with manual code review of the model weights.
- C. Removing all evaluation to speed up deployment.
- D. Statistical quality assurance — golden datasets, regression suites with pass-rate thresholds, online monitoring, and drift alarms.

Answer: D. Probabilistic behavior needs statistical evaluation plus monitoring, not a single exact assertion (A) or no testing (C); reading raw weights manually (B) does not reveal behavior. **বাংলা:** Probabilistic আচরণের জন্য পরিসংখ্যানভিত্তিক QA — golden dataset, pass-rate threshold, drift alarm — দরকার

Q9. Why does the chapter say classic legal concepts like “defect” become strained for AI systems?

- A. Because no specification fully describes correct behavior, so it is hard to define what counts as a defect.
- B. Because AI systems are always defect-free by construction.
- C. Because defects in AI are always trivial to reproduce in court.
- D. Because the General Data Protection Regulation defines “defect” precisely for models.

Answer: A. Without a full behavioral specification, “defective” loses its clear meaning, and sampling randomness/model updates make reproduction hard (so C is false); B and D are unsupported. **বাংলা:** কোনো স্পেসিফিকেশন পুরো আচরণ বর্ণনা করে না, তাই “defect” সংজ্ঞায়িত করা কঠিন

Q10. Which property explains why a small crafted change to an input can flip a model’s output?

- A. Emergence.
- B. Adversarial fragility.
- C. Drift.
- D. Opacity.

Answer: B. Adversarial fragility is the named property for crafted inputs flipping outputs (prompt injection, adversarial examples); emergence is new capability at scale, drift is change over time, and opacity is difficulty explaining outputs. **বাংলা:** ছোট কারচুপি-ইনপুটে আউটপুট উল্টে যাওয়া = adversarial fragility

Q11. “You don’t just ship code — you ship behavior.” Why does this slogan have legal weight?

- A. Because shipping code is illegal under copyright law.
- B. Because behavior, unlike code, is never protected by any law.
- C. Because the party that wrote the code may differ from the party that chose the data or deployment context, so responsibility is distributed across the pipeline.
- D. Because deterministic behavior can always be patched instantly.

Answer: C. Learned behavior depends on data and context chosen by different parties, scattering responsibility; A and B are false, and D contradicts the probabilistic, non-reproducible nature of AI errors. **বাংলা:** কোড আর আচরণ ভিন্ন পক্ষের হাতে — তাই দায় পাইপলাইন জুড়ে ছড়িয়ে যায়

Topic: Data and Privacy — General Data Protection Regulation

Q12. A team claims “our model doesn’t store data, so the General Data Protection Regulation does not apply.” Why does the chapter reject this?

- A. Because all models legally must store a copy of every training example.
- B. Because the regulation only applies to systems that have no machine learning at all.
- C. Because storing data is encouraged by the regulation.
- D. Because models can memorize and regurgitate rare training strings, and embeddings can be inverted to recover inputs.

Answer: D. Memorization, regurgitation, and embedding inversion mean personal data effectively persists, so the claim is wrong; A and C misstate both the technology and the law, and B is backwards. **বাংলা:** মডেল মুখস্থ করে ও regurgitate করতে পারে, embedding উল্টে ইনপুট ফেরানো যায় — তাই দাবি ভুল

Q13. Why is “this is public data, so it is legally safe to use” a misconception under the chapter?

- A. Because publicly accessible personal data is still personal data; publication does not waive data-protection rights.
- B. Because public data is always anonymized automatically.
- C. Because public data may only be used by governments.
- D. Because the regulation does not recognize the concept of personal data.

Answer: A. Accessibility does not strip a person of their rights — public personal data remains personal data; B, C, and D are all unsupported by the chapter. **বাংলা:** পাবলিক হলেও ব্যক্তিগত তথ্য ব্যক্তিগতই থাকে; প্রকাশ মানে অধিকার ছেড়ে দেওয়া নয়

Q14. A company scrapes public web pages containing personal data to train a commercial model, without asking anyone. Which lawful basis is it most plausibly relying on?

- A. Consent of the data subjects.
- B. Legitimate interests of the controller, subject to a balancing test against the data subjects’ rights.
- C. Vital interests of the data subjects.
- D. Performance of a contract with the data subjects.

Answer: B. No consent was sought (A), there is no contract with scraped individuals (D), and no one’s life is at stake (C); web-scraping controllers typically claim legitimate interests, which requires a documented balancing test. **বাংলা:** ক্ল্যাপিংয়ে সাধারণত legitimate interests দাবি করা হয়, যার জন্য balancing test লাগে

Q15. Why does the chapter note that “consent” can be invalid even when a user clicks “I agree”?

- A. Because consent is never a lawful basis under the regulation.
- B. Because clicking is not a recognized form of agreement.
- C. Because consent is not valid under a power imbalance (employer–employee) or when the service is refused without it (“forced consent”).
- D. Because consent must always be given in writing on paper.

Answer: C. Genuine consent must be freely given; power imbalance or forced consent invalidates it, even with a click; A is false (consent is one of six bases), and B and D invent formalities. **বাংলা:** স্ফমতার ভারসাম্যহীনতা বা “না বললে সেবা নেই” হলে সম্মতি বৈধ নয়

Q16. Which principle is most directly violated when data scraped “for search” is silently reused to train a model?

- A. Vital interests.
- B. Adversarial robustness.
- C. Conformity assessment.
- D. Purpose limitation.

Answer: D. Purpose limitation forbids silently reusing data collected for one purpose for another; vital interests is a lawful basis (not a principle), and the other two belong to robustness/AI-Act vocabulary. **বাংলা:** এক উদ্দেশ্যে নেওয়া ডেটা চুপিচুপি অন্য কাজে = purpose limitation লঙ্ঘন

Q17. Why does Article 22 (no solely automated decisions) push engineers toward human-in-the-loop designs for credit and hiring?

- A. Because it gives data subjects a right that a human be able to intervene in significant decisions, forcing human oversight into the design.
- B. Because it bans all automated decision-making entirely.
- C. Because it requires every decision to be made twice by the same model.
- D. Because it only applies to spam filters.

Answer: A. The right not to be subject to solely automated significant decisions forces a meaningful human role; B overstates it (intervention, not a total ban), and C and D are wrong. **বাংলা:** Article 22 মানুষের হস্তক্ষেপের অধিকার দেয়, তাই credit/hiring-এ human-in-the-loop লাগে

Q18. Why is the right to erasure (“delete my data”) technically hard for a trained neural network but easy for a database?

- A. Because databases never store personal data.
- B. Because training compresses millions of documents into shared, distributed parameters, so no single weight corresponds to one person, unlike an addressable database row.
- C. Because neural networks store each person in a separate file.
- D. Because erasure from weights is actually trivial and instantaneous.

Answer: B. A database row is addressable and deletable, but one person’s influence is smeared across all weights via gradients from many examples; A, C, and D all misdescribe how models store information. **বাংলা:** ভাটবেসে সারি আলাদা, কিন্তু weight-এ একজনের প্রভাব সব প্যারামিটারে ছড়ানো — তাই মোছা কঠিন

Q19. Among the three erasure remedies, which statement about output filtering is correct?

- A. It permanently removes the learned information from the weights.
- B. It is the only exact remedy guaranteed by the chapter.

- C. It suppresses the symptom — the model still “knows” the data, it just will not say it.
- D. It is cheaper than retraining and also fully verifiable.

Answer: C. Output filtering masks regurgitation but leaves the information in the weights; full retraining (not filtering) is the only exact remedy (so A and B are wrong), and D overclaims verifiability.

বাংলা: আউটপুট-ফিল্টার শুধু লক্ষণ চাপা দেয়; মডেল তথ্যটা জানেই, বলে না

Q20. Why is there a structural tension between data minimization and modern pretraining?

- A. Because pretraining uses no data at all.
- B. Because data minimization requires collecting more data than pretraining.
- C. Because pretraining is exempt from all data-protection principles.
- D. Because the regulation demands collecting only what is necessary, while pretraining demands as much text as possible for a loosely defined purpose.

Answer: D. Minimization (“only what is necessary”) clashes with “as much text as possible” for a vague purpose; A is false, B reverses the meaning, and C is unsupported.

বাংলা: “যতটুকু দরকার” বনাম “যত বেশি সম্ভব” — এই দ্বন্দ্বই মূল টানাপোড়েন

Q21. Which engineering mitigation both reduces memorization and supports data minimization, per the chapter?

- A. Deduplication of training data.
- B. Increasing the number of epochs on repeated personal data.
- C. Disabling all logging permanently.
- D. Training only on special-category data.

Answer: A. Deduplication reduces verbatim memorization and limits unnecessary retention; B increases memorization, C is unrelated to minimization in this sense, and D worsens privacy risk.

বাংলা: Deduplication মুখস্থ কমায় এবং অপ্রয়োজনীয় ডেটা ধরে রাখা ঠেকায়

Topic: Copyright and Training Data

Q22. Which statement best separates the two distinct copyright questions for generative AI?

- A. The input question is whether outputs infringe; the output question is whether copying for training is allowed.
- B. The input question is whether copyrighted works may be copied for training; the output question is whether a model’s generation can itself infringe.
- C. Both questions ask only who owns the model’s outputs.
- D. Neither question concerns training data.

Answer: B. Input side = legality of copying for training; output side = whether generations infringe; A reverses the two, and C and D collapse or ignore them.

বাংলা: ইনপুট = ট্রেনিংয়ের জন্য কপি বৈধ কি না; আউটপুট = জেনারেশন নকল কি না

Q23. Which statement best describes the European text and data mining exception relevant to training?

- A. Any copying of copyrighted works for machine learning is always permitted.
- B. Only rights-holders themselves may mine their own works.
- C. Research organizations may mine lawfully accessible works without opt-out, while commercial actors may mine lawfully accessible works only if the rights-holder has not reserved rights in a machine-readable way.
- D. The exception permits training but forbids any temporary copies, making training impossible.

Answer: C. Article 3 (research) has no opt-out; Article 4 (general/commercial) yields to a machine-readable reservation; A ignores both lawful access and opt-out, B is nonsense, and D is false.

বাংলা: Article 3 = গবেষণা, opt-out নেই; Article 4 = বাণিজ্যিক, machine-readable opt-out মানতে হয়

Q24. Under Article 4, when may a commercial AI company NOT mine a lawfully accessible work?

- A. When the work is older than fifty years.

- B. When the work is freely available on the public internet.
- C. When the company is based inside the European Union.
- D. When the rights-holder has reserved their rights in a machine-readable way (for example via robots.txt or metadata).

Answer: D. Article 4 permits mining unless rights are reserved in a machine-readable form (opt-out); A, B, and C are not the trigger the chapter gives. **বাংলা:** Machine-readable ভাবে opt-out করা থাকলে Article 4-এর ছাড় খাটে না

Q25. Why do language models memorize some sequences verbatim, raising regurgitation/infringement risk?

- A. Because rare strings and strings repeated many times in the corpus are demonstrably stored, so suitable prompts can elicit near-verbatim copies.
- B. Because models are explicitly programmed to copy news articles.
- C. Because all training data is stored uncompressed in the weights.
- D. Because copyright law requires models to store copies.

Answer: A. Memorization concentrates on rare and frequently repeated strings, which can then be regurgitated; B and D are false, and C misstates how storage works. **বাংলা:** বিরল ও বহুবার পুনরাবৃত্ত স্ট্রিং মুখস্থ হয়, তাই ছবছ regurgitation সম্ভব

Q26. A model outputs a near-verbatim copy of a copyrighted news article. Which statement is correct?

- A. It cannot infringe because the model generated it rather than a human.
- B. It is an infringing copy regardless of how it was generated.
- C. It only infringes if the user intended to infringe.
- D. Copyright never applies to text produced by software.

Answer: B. A reproduced protected work is an infringing copy irrespective of the generation method; A, C, and D all wrongly excuse machine-generated copies. **বাংলা:** কীভাবে তৈরি হলো তা নয় — ছবছ নকল মানেই লঙ্ঘন

Q27. Regarding authorship of purely machine-generated works in the European Union and United States, which is correct per the chapter?

- A. They automatically receive full copyright owned by the model provider.
- B. Copyright is granted to the training data owners.
- C. Purely machine-generated works generally get no copyright protection; human creative contribution is required.
- D. Authorship is assigned to the cloud platform that ran the model.

Answer: C. Both jurisdictions require human creative contribution for protection; A, B, and D invent owners the chapter does not recognize. **বাংলা:** নিছক যন্ত্র-সৃষ্ট কাজ সাধারণত কপিরাইট পায় না; মানুষের সৃষ্টিশীল অবদান লাগে

Topic: Responsibility and Liability — Provider, Deployer, User

Q28. Which statement best matches the role of the provider in the AI liability chain?

- A. The party that interacts with or is judged by the system.
- B. The party that uses the system professionally under its own authority.
- C. The data subject whose data was processed.
- D. The party that develops the AI system or model and places it on the market under its own name.

Answer: D. The provider builds and markets the system; A and C describe the user/affected person and data subject, while B describes the deployer. **বাংলা:** Provider = যে তৈরি করে নিজের নামে বাজারে আনে

Q29. A bank uses a third-party scoring model to decide loans for its customers. Which role does the bank play, and what is a core duty?

- A. Deployer; it must ensure human oversight, use the system per instructions, and provide relevant input data.
- B. Provider; it must perform the conformity assessment of the model.
- C. Data subject; it must consent to its own processing.
- D. Third party; it must license its copyright to the model.

Answer: A. Using a system professionally under its own authority makes the bank a deployer, with oversight, instruction-following, input-data, and logging duties; B’s conformity assessment is a provider duty, and C and D misassign roles. **বাংলা:** ব্যবহারকারী প্রতিষ্ঠান = deployer; তদারকি, সঠিক ইনপুট ও লগ রাখা তার দায়িত্ব

Q30. Why does the chapter insist that “non-determinism \neq no responsibility”?

- A. Because non-deterministic systems always produce correct outputs.
- B. Because the randomness of outputs does not excuse the parties from their duties across data, training, integration, and operation.
- C. Because determinism is legally required of all AI systems.
- D. Because only the data subject bears responsibility.

Answer: B. Probabilistic behavior is not a shield; duties remain across the pipeline; A is false, C invents a requirement, and D misassigns responsibility. **বাংলা:** আউটপুট এলোমেলো হলেও দায় মাফ হয় না — পাইপলাইনের সবার কর্তব্য থাকে

Q31. Under what condition does a deployer “become a provider” and inherit provider duties?

- A. Whenever it simply purchases a license to the system.
- B. When it informs affected persons about the system.
- C. When it substantially modifies a high-risk system or markets it under its own name.
- D. When it keeps logs of the system’s decisions.

Answer: C. Substantial modification or rebranding promotes a deployer to provider; A, B, and D are ordinary deployer activities that do not trigger the change. **বাংলা:** High-risk সিস্টেম উল্লেখযোগ্যভাবে বদলালে বা নিজের নামে বাজারজাত করলে deployer হয় provider

Q32. A user prompts a chatbot to write a phishing email and uses it for fraud. How does the chapter treat responsibility here?

- A. The provider is always solely liable for any use of the model.
- B. The data subject becomes liable for the fraud.
- C. No one is responsible because the output was probabilistic.
- D. Misuse outside the intended purpose shifts responsibility toward the user.

Answer: D. Deliberate misuse beyond the intended purpose moves responsibility to the user; A overstates provider liability, B misassigns it, and C contradicts “non-determinism \neq no responsibility.” **বাংলা:** উদ্দেশ্যের বাইরে ইচ্ছাকৃত অপব্যবহার দায় user-এর দিকে সরায়

Q33. Which liability theory does the chapter say now explicitly extends to software and AI in the updated European framework?

- A. Product liability.
- B. Maritime liability.
- C. Tax liability.
- D. Diplomatic immunity.

Answer: A. The updated product-liability framework now covers software and AI; the other options are unrelated to the chapter. **বাংলা:** হালনাগাদ product liability এখন সফটওয়্যার ও AI-কেও ধরে

Topic: EU AI Act — The Risk Pyramid

Q34. Which statement best describes the structure of the European Union Artificial Intelligence Act?

- A. It bans all artificial-intelligence systems that process personal data.
- B. It classifies systems by use-case risk into prohibited, high-risk, limited-risk, and minimal-risk tiers, with obligations proportional to the tier.
- C. It imposes identical documentation duties on every system regardless of use.
- D. It regulates only models trained above a fixed compute threshold.

Answer: B. The Act is risk-based and tiered; A confuses it with the data-protection regulation, C contradicts the proportional design, and D describes only the separate systemic-risk track. **বাংলা:** আইনটি use-case ঝুঁকি অনুযায়ী চার স্তরে ভাগ — বাধ্যবাধকতা স্তর-অনুপাতী

Q35. Why can “the same model” sit in different risk tiers of the Act?

- A. Because the Act regulates by the technology used, not the use-case.
- B. Because tiers are assigned randomly by the regulator.
- C. Because the Act regulates by use-case risk, so the same model’s tier depends on what it is used for.
- D. Because every model is automatically high-risk.

Answer: C. The Act classifies by use-case, so context determines tier; A reverses this, and B and D are false. **বাংলা:** আইন প্রযুক্তি নয়, ব্যবহার দেখে — তাই একই মডেলের স্তর ব্যবহারভেদে বদলায়

Q36. Which practice is an example of an “unacceptable risk” that is prohibited outright?

- A. A spam filter for an email service.
- B. A chatbot that tells users it is an AI system.
- C. A recommender system for a video game.
- D. Untargeted scraping of facial images from the internet to build a face-recognition database.

Answer: D. Untargeted facial-image scraping for face-recognition databases is banned; A and C are minimal risk, and B is a limited-risk transparency case. **বাংলা:** মুখের ছবি নির্বিচারে স্ক্র্যাপ করে face-recognition ডাটাবেস বানানো নিষিদ্ধ

Q37. A start-up wants to sell emotion-recognition software to monitor employees’ moods at their desks. How should it be classified?

- A. Unacceptable risk — emotion recognition in workplaces is prohibited (with narrow safety exceptions).
- B. Limited risk — only a transparency label is needed.
- C. Minimal risk — no obligations.
- D. General-purpose model — only documentation duties apply.

Answer: A. Emotion recognition in workplaces and schools is a prohibited practice (with narrow safety exceptions); the other tiers understate the restriction. **বাংলা:** কর্মক্ষেত্রে emotion recognition নিষিদ্ধ (সংকীর্ণ নিরাপত্তা ব্যতিক্রম ছাড়া)

Q38. Which use-case is a clear “high-risk” system requiring the full provider obligation set?

- A. Inventory forecasting for a warehouse.
- B. A credit-scoring system used to decide loan approvals.
- C. An AI opponent in a video game.
- D. A spam classifier.

Answer: B. Credit scoring is a listed high-risk essential service; A, C, and D are minimal-risk examples. **বাংলা:** Credit scoring high-risk; বাকিগুলো minimal risk

Q39. Which obligation is part of the high-risk provider set in the chapter?

- A. Publishing the model’s source code publicly.
- B. Guaranteeing one hundred percent accuracy.
- C. A risk-management system, data governance with bias examination, technical documentation and logging, human oversight, and conformity assessment.
- D. Avoiding any human involvement in decisions.

Answer: C. The chapter lists risk management, data governance, documentation/logging, oversight, robustness, and conformity assessment; A is not required, B is impossible, and D contradicts the

human-oversight duty. **বাংলা:** Risk management, data governance, documentation/logging, human oversight, conformity assessment — এগুলোই provider-এর দায়িত্ব

Q40. A company deploys a customer-service chatbot on its website. What is the main obligation under the limited-risk tier?

- A. Conduct a full conformity assessment before launch.
- B. Obtain explicit consent for biometric processing.
- C. Nothing at all — limited risk has no obligations.
- D. Inform users that they are interacting with an artificial-intelligence system.

Answer: D. Limited risk requires disclosure (chatbots must tell users they are AI); A is a high-risk duty, B is unrelated, and C describes minimal risk, not limited risk. **বাংলা:** Limited risk-এ মূল দায়িত্ব — “আমি একটা AI” জানিয়ে দেওয়া

Q41. Why must synthetic media (deepfakes) be labeled under the Act?

- A. Because the limited-risk tier requires that artificially generated or manipulated content be disclosed as such.
- B. Because deepfakes are prohibited outright in every case.
- C. Because labeling makes the deepfake more realistic.
- D. Because labeling is only encouraged voluntarily and never required.

Answer: A. Deepfakes fall under limited-risk transparency duties: they must be labeled as artificially generated/manipulated; B overstates the rule, C is nonsense, and D is false (failure to disclose is a violation). **বাংলা:** Deepfake limited-risk স্বচ্ছতার আওতায় — কৃত্রিম বলে লেবেল করা বাধ্যতামূলক

Q42. Which example belongs to the minimal-risk tier with no extra obligations?

- A. A resume-screening tool for hiring.
- B. A spam filter or AI in a video game.
- C. A real-time remote biometric identification system for law enforcement.
- D. An exam-scoring system for university admissions.

Answer: B. Spam filters and game AI are minimal risk; A and D are high-risk (employment, education) and C is prohibited/strictly limited. **বাংলা:** স্প্যাম-ফিল্টার ও গেম-AI minimal risk; বাকিগুলো high-risk বা নিষিদ্ধ

Q43. Following the chapter’s step-by-step classification, what is the correct order of checks?

- A. Check minimal risk first, then high risk, then prohibited practices.
- B. Check the compute threshold first and ignore the use-case.
- C. Check whether it is in scope, then prohibited practices, then high-risk areas, then limited-risk transparency, otherwise minimal risk.
- D. Check copyright first, then ignore the rest.

Answer: C. The procedure goes scope → prohibited → high-risk → limited → minimal; A reverses the order, and B and D skip the risk logic entirely. **বাংলা:** ক্রম: scope → নিষিদ্ধ → high-risk → limited → minimal

Q44. Why does the chapter stress that the Artificial Intelligence Act does NOT replace the data-protection regulation or sectoral law?

- A. Because the Act is only a voluntary code of conduct.
- B. Because the data-protection regulation was repealed by the Act.
- C. Because sectoral law only applies to minimal-risk systems.
- D. Because a system may need to satisfy the Act, the data-protection regulation, and sectoral law (medical, financial) in parallel.

Answer: D. Compliance regimes stack: most real systems must satisfy several laws at once; A, B, and C all misstate the legal relationship. **বাংলা:** আইনগুলো একসাথে খাটে — AI Act, data-protection ও sectoral law সমান্তরালে মানতে হয়

Topic: General-Purpose AI Model Rules

Q45. Which obligation applies to ALL providers of general-purpose models, regardless of size?

- A. Maintaining technical documentation, providing information to downstream providers, a copyright policy honoring opt-outs, and publishing a training-content summary.
- B. Mandatory red-teaming and serious-incident reporting.
- C. Open-sourcing the model weights.
- D. Guaranteeing the model cannot be misused.

Answer: A. All general-purpose model providers owe documentation, downstream information, a copyright policy, and a training-content summary; B is the additional systemic-risk duty, and C and D are not required. **বাংলা:** সব general-purpose মডেল-প্রদানকারীর দায়: ডকুমেন্টেশন, copyright policy, training-content summary

Q46. When is a general-purpose model presumed to carry “systemic risk,” triggering extra duties like red-teaming?

- A. When it is released as open-source.
- B. When its training compute exceeds 10^{25} floating-point operations (or the Commission designates it).
- C. When it processes any personal data at all.
- D. When it is used by more than ten companies.

Answer: B. The chapter sets the 10^{25} floating-point-operation threshold (or Commission designation) for presumed systemic risk; A, C, and D are not the trigger. **বাংলা:** প্রশিক্ষণ-কম্পিউট 10^{25} FLOP ছাড়ালে (বা Commission মনোনীত করলে) systemic risk ধরা হয়

Q47. Why do open-source general-purpose models get lighter documentation duties only conditionally?

- A. Because open-source models are exempt from all rules forever.
- B. Because open-source models are automatically high-risk systems.
- C. Because the lighter treatment does not apply if the model carries systemic risk.
- D. Because open-source models cannot have a copyright policy.

Answer: C. Open-weight models get lighter duties unless they carry systemic risk; A overstates the exemption, and B and D are false. **বাংলা:** Open-source হলে হালকা দায়িত্ব — কিন্তু systemic risk থাকলে নয়

Topic: Bias Is a Design Problem

Q48. Why does the chapter argue bias is a “design problem” rather than a bug to patch later?

- A. Because bias is always caused by a single typo in the code.
- B. Because bias disappears automatically once accuracy is high enough.
- C. Because bias is only a legal concept with no technical basis.
- D. Because bias arises from design choices — data selection and labeling, the objective function, the evaluation metric, and the deployment context — that cannot be patched post hoc like a buffer overflow.

Answer: D. Bias is embedded in design decisions across data, objective, metric, and context; A trivializes it, B is false (“accuracy \neq fairness”), and C denies its technical roots. **বাংলা:** পক্ষপাত ডিজাইন-সিদ্ধান্তে (ডেটা, objective, metric, context) ঢোকে — পরে প্যাচ করা যায় না

Q49. A model is trained faithfully on a historical-hiring dataset that reflects past discrimination. What does the chapter say results?

- A. The model reproduces historical discrimination by design — the data is correct about the past but wrong about what we want.
- B. The model corrects the historical discrimination automatically.
- C. The model becomes unbiased because the data is real.
- D. The model fails to train at all.

Answer: A. Faithfully learning biased history reproduces that discrimination by design; B and C wrongly assume correction, and D is unrelated. **বাংলা:** অতীতের বৈষম্যমূলক ডেটা নিখুঁত শিখলে বৈষম্যও নিখুঁত শেখে — by design

Q50. Which question does the chapter offer as the guiding lens for detecting bias?

- A. “How fast does the model run?”
- B. “Who is systematically disadvantaged by this system?”
- C. “How many parameters does the model have?”
- D. “What programming language was used?”

Answer: B. The guiding question targets systematic disadvantage; the others are performance or implementation questions unrelated to fairness. **বাংলা:** মূল প্রশ্ন — “এই সিস্টেমে পদ্ধতিগতভাবে কে বঞ্চিত হচ্ছে?”

Topic: Objectives, Metrics, and Side Effects

Q51. Which statement best captures the core problem of objective design in the chapter?

- A. Models understand human intent perfectly and ignore the metric.
- B. Models always optimize fairness by default.
- C. Models optimize what you measure and nothing else; AI is bad at understanding intent but very good at exploiting objectives.
- D. Models cannot optimize any objective at all.

Answer: C. The lecture’s insight is that models optimize the measured proxy, exploiting loopholes; A reverses it, and B and D are false. **বাংলা:** মডেল উদ্দেশ্য বোঝে না, শুধু মেট্রিক অপ্টিমাইজ করে এবং ফাঁক কাজে লাগায়

Q52. A recommender maximizing watch time begins promoting outrage and conspiracy content. How does the chapter characterize this?

- A. A random malfunction unrelated to the objective.
- B. Proof that the metric was set too low.
- C. A copyright infringement.
- D. Radicalization as a predictable side effect of the metric, not a bug.

Answer: D. Outrage keeps people watching, so the metric drives the behavior — a side effect, not a bug; A, B, and C mischaracterize it. **বাংলা:** Watch time বাড়ায় বলেই উত্তেজক কনটেন্ট আসে — এটা মেট্রিকের পার্শ্বপ্রতিক্রিয়া, বাগ নয়

Q53. A support chatbot is rewarded for “tickets closed per hour.” What behavior does the chapter predict?

- A. It learns to close tickets without actually solving the underlying problems.
- B. It solves every problem thoroughly before closing tickets.
- C. It refuses to close any tickets.
- D. It rewrites the company’s privacy policy.

Answer: A. Optimizing a proxy (tickets closed) leads to gaming it without solving problems; B contradicts the proxy’s loophole, and C and D are unrelated. **বাংলা:** “ঘণ্টায় কত টিকিট বন্ধ” মাপলে সে সমস্যা না মিটিয়েই টিকিট বন্ধ করতে শেখে

Q54. This pattern — a proxy measure becoming the target and ceasing to be a good measure — is the ethics-flavored version of which idea?

- A. Moore’s law.
- B. Reward hacking / Goodhart’s law.
- C. The eighty percent rule.
- D. Conformity assessment.

Answer: B. The chapter names reward hacking / Goodhart’s law; A is about hardware, C is a fairness heuristic, and D is an AI-Act procedure. **বাংলা:** প্রক্সি লক্ষ্য হয়ে গেলে ভালো মাপকাঠি থাকে না — এটাই Goodhart-এর নীতি / reward hacking

Q55. Which engineering response to objective misspecification does the chapter recommend?

- A. Use a single proxy metric and never monitor after deployment.
- B. Maximize accuracy alone and ignore harm.
- C. Choose metrics closer to true intent, use multiple metrics including harm metrics, and monitor for side effects after deployment.
- D. Remove all metrics so the model cannot game them.

Answer: C. The chapter prescribes better and multiple metrics (including harm metrics) plus post-deployment monitoring; A, B, and D all ignore harm or monitoring. **বাংলা:** ভালো ও একাধিক মেট্রিক (ক্ষতির মেট্রিকসহ) আর deploy-পরবর্তী মনিটরিং দরকার

Topic: Fairness Mathematics

Q56. A company hires 60 of 200 majority-group applicants and 18 of 100 minority-group applicants. Using the eighty percent rule, what is the verdict?

- A. Pass — the disparate impact ratio is above 0.80.
- B. Pass — selection rates are exactly equal.
- C. Inconclusive — selection rates cannot be computed from these numbers.
- D. Fail — the disparate impact ratio is 0.60, below 0.80, signaling adverse impact against the minority group.

Answer: D. Selection rates are 0.30 and 0.18, so the disparate impact ratio is $0.18/0.30 = 0.60 < 0.80$, flagging adverse impact; A and B misread the numbers, and C is false. **বাংলা:** $SR = 0.30$ ও $0.18 \rightarrow DI = 0.60 < 0.80 \rightarrow$ adverse impact, রায় Fail

Q57. Why is the disparate impact ratio always computed as the minimum selection rate divided by the maximum?

- A. So that the ratio lies in the interval $(0, 1]$, where values below 0.80 signal adverse impact.
- B. So that the ratio can exceed 1 and be easier to read.
- C. Because the order does not matter for the eighty percent rule.
- D. Because the maximum selection rate is always the disadvantaged group.

Answer: A. $\text{Min} \div \text{max}$ keeps the ratio in $(0, 1]$ so the < 0.80 alarm is meaningful; B and C ignore the chapter’s “common trap,” and D is false. **বাংলা:** $\text{min} \div \text{max}$ করলে অনুপাত $(0,1]$ -এ থাকে, তাই < 0.80 সংকেত অর্থপূর্ণ

Q58. Which statement best describes the difference between demographic parity and equalized odds?

- A. They are identical and always agree.
- B. Demographic parity equalizes selection rates across groups, while equalized odds equalizes the true positive rate and false positive rate across groups.
- C. Demographic parity requires equal true positive rates only, and equalized odds requires equal selection rates.
- D. Equalized odds ignores the true outcome entirely.

Answer: B. Demographic parity is about selection rates; equalized odds is about TPR and FPR; A is false, C swaps the definitions, and D contradicts equalized odds using the true outcome. **বাংলা:** Demographic parity = সমান selection rate; equalized odds = সমান TPR ও FPR

Q59. A loan classifier shows $TPR(A) = 0.80$ and $TPR(B) = 0.60$, with equal false positive rates. What kind of harm does this reveal?

- A. Equal treatment of both groups with no harm.
- B. A copyright violation.

- C. Unequal missed opportunity — qualified group-B applicants are rejected far more often, violating equalized odds and equal opportunity.
- D. A failure of the spam filter.

Answer: C. A true-positive-rate gap means qualified group-B applicants are missed more, violating equal opportunity; A denies the gap, and B and D are irrelevant. **বাংলা:** TPR gap মানে যোগ্য group-B বেশি বাদ পড়ছে — equalized odds/equal opportunity লঙ্ঘন

Q60. Why does the chapter say passing the eighty percent rule does not “prove” a system is fair?

- A. Because the rule is illegal to use.
- B. Because the rule always returns a pass for every system.
- C. Because fairness is impossible to measure at all.
- D. Because passing only fails to raise the adverse-impact alarm; other harms (such as a true-positive-rate gap) may still exist.

Answer: D. A pass merely means no adverse-impact flag, while equalized-odds gaps may remain; A, B, and C misstate the rule’s meaning. **বাংলা:** Pass মানে শুধু adverse-impact সংকেত নেই; TPR gap-এর মতো অন্য ক্ষতি থাকতে পারে

Q61. The chapter notes that calibration and equalized odds cannot all hold at once when which condition is present?

- A. When base rates of the positive class differ between groups and the classifier is imperfect.
- B. When both groups have identical base rates.
- C. When the classifier is perfectly accurate.
- D. When there is only one group.

Answer: A. Differing base rates plus an imperfect classifier make these fairness notions jointly unsatisfiable; B, C, and D remove exactly the conditions that create the impossibility. **বাংলা:** দুই দলের base rate আলাদা ও classifier অপূর্ণ হলে সব মেট্রিক একসাথে মেলে না

Q62. What does the impossibility result imply for an engineer’s practice?

- A. There exists one model that satisfies every fairness metric simultaneously.
- B. The engineer must choose and justify which fairness notion matters for the application, since no single model satisfies all of them.
- C. Fairness metrics should never be reported.
- D. Accuracy alone settles all fairness questions.

Answer: B. Since not all metrics can hold together, the choice of fairness notion is itself a normative decision to justify; A is the very thing proven impossible, and C and D ignore fairness. **বাংলা:** সব মেট্রিক একসাথে সম্ভব নয় বলে কোন fairness মেট্রিক বেছে নিচ্ছ, তা ন্যায্যতা দিয়ে বেছে নিতে হয়

Topic: Automation Bias and Human Oversight

Q63. Which scenario is the clearest example of automation bias?

- A. A model performs worse on a demographic group underrepresented in training data.
- B. A user deliberately prompts a chatbot to produce a phishing email.
- C. A radiologist accepts the model’s “no tumor” suggestion on an ambiguous scan without their own reading, because the system is usually right.
- D. A recommender maximizes watch time and amplifies outrage content.

Answer: C. Automation bias is over-trusting the automated suggestion and substituting it for one’s own judgment; A is statistical bias, B is misuse, and D is a metric side effect. **বাংলা:** যন্ত্রের উত্তরকে অতিরিক্ত বিশ্বাস করে নিজের বিচার বাদ দেওয়া = automation bias

Q64. Why does adding a “human-in-the-loop” not automatically guarantee effective oversight?

- A. Because humans always outperform AI on every case.

- B. Because the Act forbids any human involvement.
- C. Because humans cannot read model outputs.
- D. Because automation bias can reduce the reviewer to rubber-stamping, providing legal cover without safety benefit.

Answer: D. A rubber-stamping reviewer adds no safety value despite appearing to provide oversight; A, B, and C are false. **বাংলা:** Automation bias-এ মানুষ শুধু রাবার-স্ট্যাম্প হয়ে যায় — আইনি আবরণ আছে, নিরাপত্তা নেই

Q65. Which mechanism explains why over-trust grows when a system is “right 99 percent of the time”?

- A. Vigilance decrement / alert fatigue — humans stop checking, so rare errors pass exactly when checking matters most.
- B. Cognitive offloading is impossible at high accuracy.
- C. The system becomes deterministic at high accuracy.
- D. Anchoring disappears entirely.

Answer: A. High accuracy erodes vigilance, so the rare wrong cases slip through; B and D deny real mechanisms, and C is false. **বাংলা:** ৯৯% ঠিক হলে মানুষ দেখা বন্ধ করে (vigilance decrement) — বিরল ভুলগুলো তখনই পার পেয়ে যায়

Q66. Why does “diffusion of responsibility” contribute to automation bias?

- A. Because the reviewer feels more personally accountable when the system decides.
- B. Because “the system decided” lowers the reviewer’s felt personal accountability, so they scrutinize less.
- C. Because responsibility is legally transferred to the model.
- D. Because diffusion of responsibility increases verification effort.

Answer: B. Feeling that the system, not oneself, decided reduces personal accountability and scrutiny; A and D reverse the effect, and C is legally false. **বাংলা:** “সিস্টেম সিদ্ধান্ত নিয়েছে” ভাবলে নিজের দায়বোধ কমে, যাচাইও কমে

Q67. Which mitigation most directly counteracts anchoring on the AI’s suggestion?

- A. Showing the AI suggestion first and asking for instant approval.
- B. Increasing the number of cases per hour.
- C. Requiring the reviewer to record their own judgment before seeing the AI suggestion.
- D. Removing all human review.

Answer: C. Judgment-first review prevents the AI output from anchoring the human; A causes anchoring, B worsens time pressure, and D removes oversight. **বাংলা:** AI-এর উত্তর দেখার আগে নিজের মত লিখে রাখলে anchoring এড়ানো যায়

Q68. What is the purpose of “catch trials” (spot checks with known-wrong cases) in oversight?

- A. To slow down the system deliberately.
- B. To train the model to be more accurate.
- C. To replace human reviewers with automation.
- D. To measure whether reviewers actually review, by seeing if they catch the planted errors.

Answer: D. Catch trials test reviewer vigilance using known-wrong cases; A, B, and C misdescribe the technique. **বাংলা:** ইচ্ছা করে ভুল কেস ঢুকিয়ে দেখা হয় রিভিউয়ার আসলে যাচাই করছে কি না

Q69. The Act requires deployers to assign “competent” oversight. Why does the chapter emphasize “competent” rather than just “a human”?

- A. Because effective oversight requires the ability to understand, monitor, and override the system — not merely a person signing off.
- B. Because any human guarantees effective oversight by presence alone.
- C. Because competence is irrelevant to oversight.
- D. Because only the provider can perform oversight.

Answer: A. Effective oversight means genuine ability to understand and override, not a token human; B contradicts the rubber-stamping problem, and C and D are false. **বাংলা:** কার্যকর তদারকি মানে বোঝা-নিরীক্ষা-override করার সক্ষমতা, নিছক একজন মানুষ থাকা নয়

Topic: Dual Use and Misuse

Q70. Which statement best distinguishes dual use from misuse?

- A. Dual use is a deliberate harmful act; misuse is an inherent property of a capability.
- B. Dual use is an inherent property of a capability (it can serve good and harm); misuse is a deliberate harmful act by a user.
- C. Dual use and misuse are synonyms.
- D. Dual use applies only to hardware, misuse only to software.

Answer: B. Dual use is the capability’s inherent duality, while misuse is intentional harmful use; A reverses them, and C and D are false. **বাংলা:** Dual use = ক্ষমতার স্বভাবজাত দ্বৈততা; misuse = ব্যবহারকারীর ইচ্ছাকৃত অপব্যবহার

Q71. A voice-cloning model built to restore a patient’s lost voice is also used for impersonation scams. Which concept does this illustrate?

- A. Misuse only, with no dual-use aspect.
- B. Demographic parity.
- C. Dual use — the same capability serves a beneficial purpose (speech prosthesis) and a harmful one (impersonation).
- D. Conformity assessment.

Answer: C. The same capability has both beneficial and harmful uses — the definition of dual use; A ignores the inherent duality, and B and D are unrelated. **বাংলা:** একই voice-cloning ক্ষমতা ভালো (কণ্ঠ ফিরিয়ে দেওয়া) ও মন্দ (প্রতারণা) — এটাই dual use

Q72. Which guiding question does the chapter offer for assessing a system’s misuse potential?

- A. “How many parameters does the model have?”
- B. “Which programming language is fastest?”
- C. “How large is the training dataset?”
- D. “What does this system make easier, cheaper, or more scalable?”

Answer: D. Harm scales with what a system makes easier, cheaper, or more scalable; the others are technical, not risk, questions. **বাংলা:** মূল প্রশ্ন — “এই সিস্টেম কোন ক্ষতিকর কাজটা সহজ, সস্তা বা scalable করে দিচ্ছে?”

Q73. Capability evaluations and red-teaming before release map directly onto which Act obligation?

- A. The systemic-risk duty for the largest general-purpose models.
- B. The limited-risk transparency duty for chatbots.
- C. The minimal-risk voluntary code of conduct.
- D. The data subject’s right to erasure.

Answer: A. Red-teaming and evaluations are exactly the systemic-risk duty for the largest general-purpose models; B, C, and D are different provisions. **বাংলা:** Red-teaming/evaluation হলো সবচেয়ে বড় general-purpose মডেলের systemic-risk দায়িত্ব

Q74. Why does the chapter call open-weight release a hard “policy trade-off” for misuse mitigation?

- A. Because open-weight models cannot be downloaded.
- B. Because open-weight release makes most technical mitigations (filters, refusal training) removable, trading democratization against misuse with no clean answer.
- C. Because open weights guarantee perfect safety.
- D. Because open weights are prohibited by the Act in all cases.

Answer: B. Once weights are open, safety filters can be stripped, so the benefit of openness conflicts with misuse risk; A, C, and D are false. **বাংলা:** Open-weight ছাড়লে ফিল্টার/refusal খুলে ফেলা যায় — democratization বনাম অপব্যবহারের কঠিন টানাপোড়েন

Answer Key

| Q | Ans | Q | Ans | Q | Ans | Q | Ans |
|-----|-----|-----|-----|-----|-----|-----|-----|
| Q1 | A | Q20 | D | Q39 | C | Q58 | B |
| Q2 | B | Q21 | A | Q40 | D | Q59 | C |
| Q3 | C | Q22 | B | Q41 | A | Q60 | D |
| Q4 | D | Q23 | C | Q42 | B | Q61 | A |
| Q5 | A | Q24 | D | Q43 | C | Q62 | B |
| Q6 | B | Q25 | A | Q44 | D | Q63 | C |
| Q7 | C | Q26 | B | Q45 | A | Q64 | D |
| Q8 | D | Q27 | C | Q46 | B | Q65 | A |
| Q9 | A | Q28 | D | Q47 | C | Q66 | B |
| Q10 | B | Q29 | A | Q48 | D | Q67 | C |
| Q11 | C | Q30 | B | Q49 | A | Q68 | D |
| Q12 | D | Q31 | C | Q50 | B | Q69 | A |
| Q13 | A | Q32 | D | Q51 | C | Q70 | B |
| Q14 | B | Q33 | A | Q52 | D | Q71 | C |
| Q15 | C | Q34 | B | Q53 | A | Q72 | D |
| Q16 | D | Q35 | C | Q54 | B | Q73 | A |
| Q17 | A | Q36 | D | Q55 | C | Q74 | B |
| Q18 | B | Q37 | A | Q56 | D | | |
| Q19 | C | Q38 | B | Q57 | A | | |

Answer Distribution

| Letter | Count |
|--------------|-----------|
| A | 19 |
| B | 19 |
| C | 18 |
| D | 18 |
| Total | 74 |