第一大題 (40%)：單一選擇題 (每題二分，共計四十分)，請於【選擇題作答區】內作答。

1. A renowned investment bank in Frankfurt has designed a unique assessment exercise for its graduate recruitment program. The exercise, called the "Strategic Selection Protocol," is conducted as follows:

Four identical envelopes are presented to the candidate. One envelope contains an invitation to join the prestigious Mergers & Acquisitions team with a guaranteed first-year bonus of €150,000. The other three envelopes each contain an assignment to the Operations department with a standard bonus structure. The candidate has no initial information about which envelope contains the M&A invitation. The assessment proceeds in stages:

Stage 1: The candidate selects one envelope but does not open it.

Stage 2: The managing director, who knows the contents of all envelopes, deliberately opens two of the three remaining envelopes, always revealing two Operations assignments. If the candidate's initial selection contains the M&A invitation, the director randomly chooses which two of the three Operations envelopes to open. If the candidate's initial selection contains an Operations assignment, the director opens two of the remaining Operations envelopes (leaving the M&A envelope and one Operations envelope unopened).

Stage 3: The candidate is offered the opportunity to either retain their original selection or switch to the other unopened envelope.

Mr. Weber, a candidate from Munich, initially selects Envelope A. The managing director then opens Envelopes C and D, both revealing Operations assignments. Mr. Weber must now decide whether to keep Envelope A or switch to Envelope B.

If Mr. Weber switches to Envelope B, what is the probability that he receives the M&A invitation?

(A) 1/4
(B) 1/2
(C) 2/3
(D) 3/4
(E) 1/3

2. A multinational consulting firm based in Mumbai has designed an innovative case interview format to assess candidates' persistence and analytical capabilities. The interview room contains a complex decision-simulation system with four pathways, each represented by a door of identical appearance.

The pathways are engineered as follows:

Door A leads to immediate success (job offer) after a 15-minute problem-solving exercise.

Door B leads to immediate success (job offer) after a 25-minute problem-solving exercise.

Door C leads to a 40-minute exercise that returns the candidate to the starting room with no indication of having returned.

Door D leads to a 60-minute exercise that returns the candidate to the starting room with no indication of having returned.

Upon each return to the starting room, the doors are randomized in appearance. The candidate has no memory of which doors were previously attempted and selects each door with equal probability 1/4, independently of all previous selections.

Ms. Sharma enters the interview room at exactly 10:00 AM.

What is the expected total time (in minutes) for Ms. Sharma to receive a job offer?

(A) 35 minutes
(B) 70 minutes
(C) 90 minutes
(D) 120 minutes
(E) 140 minutes

3. A leading intellectual property law firm headquartered in Boston specializes in patent litigation for technology companies. The firm's managing partners are developing a revenue forecasting model for the upcoming fiscal year based on their historical case data.

見背面

The firm's analytics team has established the following stochastic model:

Case Acquisition: The number of new patent litigation cases assigned to a randomly selected senior partner during a fiscal year follows a Poisson distribution with parameter $\lambda = 6.0$ cases.

Case Success Rate: Each case, independently of all other cases, has a probability $p = 0.35$ of resulting in a "favorable verdict"—defined as either a court ruling in favor of the client or a settlement exceeding the client's initial demands.

The firm's CFO is preparing financial projections for the partnership meeting. He explains to a junior associate: "Let N denote the number of cases assigned to a partner, and let X denote the number of favorable verdicts achieved by that partner. We need to understand both the expected number of favorable verdicts and the variability around that expectation."

The junior associate responds: "I calculated $E[X] = \lambda p = 2.1$. For the variance, since X|N follows a Binomial distribution, I used the law of total variance. My calculation shows $Var(X) = \lambda p(1-p) = 1.365$."

The CFO replies: "Your expected value is correct, but check your variance calculation again. You're missing something important."

Which of the following correctly states E[X] and Var(X)?

(A) $E[X] = 2.10$ and $Var(X) = 1.365$

(B) $E[X] = 2.10$ and $Var(X) = 2.10$

(C) $E[X] = 2.10$ and $Var(X) = 2.835$

(D) $E[X] = 3.90$ and $Var(X) = 2.10$

(E) $E[X] = 2.10$ and $Var(X) = 1.735$

4. A large health insurance company headquartered in Amsterdam launched a voluntary "Digital Health Engagement Program" (DHEP) designed to encourage policyholders to adopt healthier lifestyles through a smartphone application. The program offers personalized fitness recommendations, nutrition tracking, and virtual health coaching.

The company's actuarial division conducted a two-year retrospective study to evaluate the program's impact on healthcare costs. At the program's launch, all 120,000 eligible policyholders (ages 30-55, no pre-existing chronic conditions) were invited to participate. Enrollment was entirely voluntary.

28,000 policyholders enrolled in and actively used the DHEP.

92,000 policyholders did not enroll in the program.

The actuarial team measured each policyholder's "Annual Claims Ratio" (ACR)—the ratio of actual medical claims to expected claims based on demographic factors:

| Group | Mean ACR |
|---|---|
| DHEP Participants | 0.72 |
| Non-Participants | 0.94 |

The Chief Actuarial Officer initially reports to the executive board: "The DHEP program reduces the claims ratio by 0.22, representing a 23% reduction in healthcare costs. This demonstrates exceptional ROI for our wellness investment."

However, further analysis of pre-program characteristics reveals that participants had significantly lower baseline BMI (24.2 vs. 28.7), higher gym membership rates (62% vs. 18%), lower smoking rates (8% vs. 24%), and higher household income ($125,000 vs. $78,000).

Which of the following BEST explains why the observed 0.22 reduction in claims ratio cannot be interpreted as the causal effect of the DHEP program?

(A) Measurement error in the claims ratio creates systematic bias that inflates the apparent program effect.

(B) Confounding exists because policyholders who chose to enroll were already healthier and more health-conscious, and would have had lower claims regardless of program participation.

(C) The Hawthorne effect causes participants to temporarily change behavior due to being observed, not due to the program itself.

接次頁

(D) Survivorship bias occurs because policyholders who dropped out of the program are excluded from the analysis.
(E) Ecological fallacy occurs because group-level statistics cannot be applied to individual policyholders.

5. A prestigious business school in Chicago offers an optional "Executive Mentorship Program" (EMP) connecting current MBA students with alumni executives. The career services office conducted a study to evaluate the program's impact on post-graduation salary outcomes.
Observed Outcomes (Base Salary, in thousands USD):

| Group | Mean Salary |
|---|---|
| EMP Participants (180 students) | $168,000 |
| Non-Participants (420 students) | $142,000 |

Natural Experiment: Due to a scheduling system malfunction during the registration period, a random subset of 75 students who attempted to register for EMP were unable to complete enrollment before the deadline. This technical failure was completely random and unrelated to student characteristics.
Outcome Data for System-Affected Students (None received mentorship):

| Registration Attempt | Mean Salary |
|---|---|
| Tried to Register (25 students) | $156,000 |
| Did NOT Try to Register (50 students) | $138,000 |

Using the decomposition: Observed Difference = True Causal Effect (ATT) + Selection Bias
Using the natural experiment data, what is the magnitude of the selection bias in the original observational comparison?
(A) $12,000
(B) $14,000
(C) $18,000
(D) $26,000
(E) $30,000

6. A national retail chain headquartered in Toronto is evaluating the effectiveness of a new employee training program designed to improve customer service quality. The program, called "Service Excellence Initiative" (SEI), consists of intensive workshops on customer engagement, conflict resolution, and upselling techniques.
The company's HR Analytics team conducted a six-month observational study across 200 retail locations. The primary outcome measure is the "Customer Satisfaction Index" (CSI)—a composite score from 0 to 100 based on customer feedback surveys.
Observed Adoption and Outcomes:

| Group | Mean CSI |
|---|---|
| SEI Participants (500 employees) | 84.5 |
| Non-Participants (1,500 employees) | 71.5 |

Instrumental Variable: During the study period, 80 of the 200 retail locations received a corporate mandate requiring store managers to allocate dedicated training hours for staff development. This mandate was issued based on regional distribution center logistics schedules established two years prior—completely independent of individual employee or store characteristics.

見背面

| | WITH Mandate | WITHOUT Mandate |
|---|---|---|
| SEI participation rate | 45% | 15% |
| Mean CSI | 78.9 | 72.9 |

The LATE formula is: LATE = (Difference in outcomes by instrument) / (Difference in participation rates by instrument)

Based on the instrumental variable analysis, what is the estimated true causal effect (LATE) of SEI participation on Customer Satisfaction Index?

(A) 6.0 points
(B) 13.0 points
(C) 15.0 points
(D) 20.0 points
(E) 24.0 points

7. A leading pharmaceutical corporation headquartered in Basel, Switzerland, is conducting a critical Phase III clinical trial for a novel oncology drug designated as XR-7721. The drug's efficacy is measured by the "Tumor Response Index" (TRI), a continuous metric ranging from 0 to 1, where higher values indicate greater tumor shrinkage. Regulatory agencies require that the drug demonstrate a median TRI of at least 0.5 and a mean TRI of at least 0.45 to receive approval for market authorization.

The corporation's biostatistics division has developed a probability model based on preliminary Phase II data. They propose that the TRI for a randomly selected patient follows a continuous distribution with probability density function:

$$f(x) = c(1 + kx^2) \quad \text{for } 0 \le x \le 1, \quad \text{and 0 otherwise}$$

where c and k are constants. Extensive pharmacokinetic modeling suggests that the parameter k equals 2 for the target patient population.

The biostatistics team presents the following assertions to the executive committee ahead of the regulatory submission:

(I) For the density function to be valid when $k = 2$, the normalizing constant c must equal 3/5.

(II) Under this model, the probability that a randomly selected patient exhibits a TRI below 0.5 is 11/40.

(III) The expected value of the TRI under this model is 3/5, which exceeds the regulatory threshold of 0.45.

(IV) The cumulative distribution function evaluated at $x = 0.5$ can be expressed as $F(0.5) = (3/5)(1/2 + 1/12)$.

(V) If the regulatory agency instead required that $P(TRI > 0.75) \ge 0.50$ for approval, the drug would fail to meet this criterion under the proposed model.

Which of the above assertions are CORRECT?

(A) I, II, and IV only
(B) I, III, and V only
(C) I, III, IV, and V only
(D) II, III, and IV only
(E) I, II, III, and V only

8. A multinational management consulting firm based in London operates a performance-based compensation system for its associates. Each associate's quarterly bonus is determined by a "Client Impact Score" (CIS), a discrete random variable that reflects the measurable value delivered to clients during the quarter. The firm's compensation committee has analyzed historical performance data and established that the CIS for a randomly selected associate follows the probability distribution shown below:

| CIS (X) | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| P(X=x) | 0.10 | $p$ | 0.30 | $q$ | 0.15 |

接次頁

The committee's actuarial analysis indicates that E[X] = 3.05.

The firm employs a tiered bonus structure across its three regional offices. Let X denote the CIS for a randomly selected associate. The quarterly bonus (in thousands of GBP) is calculated as follows:

London office: $B_L = 8X + 5$

Singapore office: $B_S = 6X + 12$

New York office: $B_N = 10X - 2$

The Chief Financial Officer is evaluating whether to standardize the bonus structure globally. After solving for the unknown parameters p and q, the CFO asks an analyst to compare the three compensation schemes.

Which of the following statements is TRUE?

(A) The New York office has the highest expected bonus, and the London office has the highest variance in bonus payments.

(B) The Singapore office has the highest expected bonus, and the New York office has the highest variance in bonus payments.

(C) All three offices have the same expected bonus, but the New York office has the highest variance in bonus payments.

(D) The London office has the highest expected bonus, and the Singapore office has the lowest variance in bonus payments.

(E) The Singapore office has the highest expected bonus, and the London office has the lowest variance in bonus payments.

9. A prestigious private equity firm in Hong Kong has developed an unconventional final-round interview process designed to assess candidates' decision-making capabilities under uncertainty. The process, known internally as the "Strategic Pivot Challenge," operates as follows:

Three identical briefcases are presented to the candidate. One briefcase contains a binding offer letter for a senior associate position with a signing bonus of HK$2,000,000. The other two briefcases each contain a polite rejection letter. The candidate has no initial information regarding which briefcase contains the offer.

The interview proceeds in stages:

Stage 1: The candidate selects one briefcase but does not open it.

Stage 2: The managing director, who knows the contents of all briefcases, deliberately opens one of the two remaining briefcases, always revealing a rejection letter. Crucially, if the candidate's initial selection contains the offer, the director randomly chooses which of the two rejection-containing briefcases to open with equal probability. If the candidate's initial selection contains a rejection, the director opens the only other briefcase that contains a rejection.

Stage 3: The candidate is offered the opportunity to either retain their original selection or switch to the other unopened briefcase.

A candidate named Ms. Chen initially selects Briefcase A. The managing director then opens Briefcase B, revealing a rejection letter.

However, this year the firm introduces a modification to assess probabilistic reasoning at a deeper level. Before Ms. Chen decides whether to switch, she is informed that a junior analyst—who was unaware of the briefcase contents—entered the room during Stage 2 and observed only that Briefcase B was opened and contained a rejection letter. The junior analyst did not observe Ms. Chen's initial selection.

From the perspective of the junior analyst, what is the probability that Briefcase C contains the offer letter?

(A) 1/3

(B) 1/2

(C) 2/3

(D) 3/4

(E) The probability cannot be determined without knowing Ms. Chen's initial selection.

10. A global management consulting firm has designed an elaborate assessment center for its final-round recruitment process at its headquarters in Boston. The centerpiece of this assessment is an architectural maze

見背面

constructed in the basement level of the building, intended to evaluate candidates' resilience, decision-making under uncertainty, and psychological composure.

The maze operates as follows: A candidate enters and immediately encounters a junction with three visually indistinguishable corridors—labeled internally by the facility engineers as Corridor X, Corridor Y, and Corridor Z, though candidates cannot differentiate among them. The corridors are engineered such that:

- Corridor X leads directly to the exit after exactly 20 minutes of walking.
- Corridor Y leads directly to the exit after exactly 35 minutes of walking.
- Corridor Z is a deceptive loop that returns the candidate to the original junction after exactly 110 minutes of walking, with no indication that they have returned to their starting point.

The junction room is designed with rotating architecture such that, upon returning via Corridor Z, the three corridors appear in a completely randomized configuration. The candidate retains no memory of which corridor was previously attempted, as all corridors are identical in appearance. At each visit to the junction, the candidate selects one of the three corridors with equal probability 1/3, independently of all previous selections.

Mr. Rodriguez, an MBA candidate from São Paulo, enters the maze at exactly 9:00 AM.

What is the expected total time (in minutes) for Mr. Rodriguez to exit the maze?

(A) 55 minutes
(B) 65 minutes
(C) 80 minutes
(D) 82.5 minutes
(E) 110 minutes

11. A premier venture capital firm based in Menlo Park, California, has developed a sophisticated quantitative framework for modeling its investment pipeline. The firm's analytics team, composed of former quantitative researchers from leading hedge funds and technology companies, has been tasked with forecasting the expected number of successful exits for the upcoming fiscal year to guide capital reserve planning and limited partner communications.

The firm's historical data, accumulated over two decades of investment activity across multiple fund vintages, supports the following stochastic model:

***Deal Flow Model:*** The number of qualified investment opportunities that reach the final due diligence stage for a randomly selected partner during a fiscal year follows a Poisson distribution with parameter $\lambda = 4.5$ deals.

***Conversion Rate:*** Each deal that reaches final due diligence, independently of all other deals, has a probability $p = 0.20$ of ultimately becoming a "successful exit"—defined as an IPO or acquisition that returns at least 3x the invested capital within seven years.

The firm's Managing Director is preparing for the annual Limited Partners Advisory Committee meeting. She convenes the analytics team to address the following question:

"Consider a randomly selected partner from our firm. Let N denote the number of deals reaching final due diligence for this partner during the fiscal year, and let X denote the number of these deals that ultimately become successful exits. Our LP from the Yale Endowment—herself a former statistician—will certainly ask about the variance of X, not just its expected value."

A junior analyst, recently graduated from a top MBA program, quickly calculates: "The expected number of successful exits per partner is simply $\lambda p = 0.9$. For the variance, I applied the law of total variance and obtained $\text{Var}(X) = E[\text{Var}(X|N)] = \lambda p(1-p) = 0.72$."

The Chief Risk Officer shakes his head. "You've forgotten a term. The law of total variance has two components."

Which of the following correctly states E[X] and Var(X)?

(A) $E[X] = 0.90$ and $\text{Var}(X) = 0.72$
(B) $E[X] = 0.90$ and $\text{Var}(X) = 0.90$
(C) $E[X] = 0.90$ and $\text{Var}(X) = 1.08$
(D) $E[X] = 4.50$ and $\text{Var}(X) = 0.90$

接次頁

(E) E[X] = 3.60 and Var(X) = 0.72

12. A major retail banking corporation headquartered in Singapore is evaluating the effectiveness of a new financial literacy program designed to increase customers' adoption of long-term investment products. The program, branded "WealthWise Academy," consists of a series of online modules covering topics such as compound interest, portfolio diversification, risk-adjusted returns, and retirement planning.

The bank's Customer Analytics division conducted a study over a twelve-month period. At the beginning of the study, the program was made available to all 48,000 retail banking customers in the target demographic (ages 28-45, annual income above SGD 80,000). Customers could voluntarily choose whether to enroll. By the end of the study period:

- 12,000 customers enrolled in and completed the WealthWise Academy program.
- 36,000 customers did not enroll in the program.

The analytics team measured each customer's "Investment Portfolio Growth Rate" (IPGR)—the percentage increase in total invested assets during the twelve-month period:

| **Group** | **Mean IPGR** |
|---|---|
| Program Participants | 18.7% |
| Non-Participants | 11.2% |

The Chief Marketing Officer initially reports: "The WealthWise Academy program caused a 7.5 percentage point increase in customers' investment portfolio growth."

The Chief Risk Officer reviews pre-program characteristics and finds that participants had significantly higher prior-year IPGR (14.2% vs. 8.6%), higher financial literacy scores (72.4 vs. 51.8), higher proportion with graduate degrees (68% vs. 34%), and more prior investment accounts (3.8 vs. 1.9).

Which of the following BEST explains why the observed 7.5 percentage point difference cannot be interpreted as the causal effect of the WealthWise Academy program?

(A) The sample sizes are unequal between the two groups, which biases the comparison of means.

(B) Selection bias arises because customers who chose to enroll likely differ systematically from non-enrollees in unobserved ways that independently affect investment behavior, violating the ignorability assumption.

(C) The study lacks statistical power because the standard deviations are too large relative to the observed difference.

(D) The comparison is invalid because the IPGR is influenced by macroeconomic factors that affected all customers equally.

(E) Regression to the mean explains the entire observed difference.

13. A leading technology conglomerate based in Shenzhen, China, implemented a voluntary "Executive Leadership Acceleration Program" (ELAP) for its mid-level managers. The primary outcome measure is the "Performance Rating Improvement" (PRI)—the change in annual performance rating (on a 100-point scale) from the year before program availability to two years after.

Observed Outcomes:

| Group | Mean PRI |
|---|---|
| ELAP Participants (400 managers) | +12.6 points |
| Non-Participants (1,600 managers) | +5.4 points |

Natural Experiment: Due to an administrative error during enrollment, a randomly selected subset of 300 managers experienced a system glitch that prevented them from accessing the enrollment portal. This glitch was completely random and unrelated to any manager characteristics.

Outcome Data for Glitch-Affected Managers (None received treatment):

見背面

| Intention | Mean PRI |
|---|---|
| Intended to Enroll (60 managers) | +10.8 points |
| Did NOT Intend to Enroll (240 managers) | +4.2 points |

The decomposition framework is: Observed Difference = True Causal Effect (ATT) + Selection Bias.
Using the natural experiment data to estimate the counterfactual, what is the magnitude of the selection bias in the original observational comparison?
(A) 1.8 points
(B) 4.2 points
(C) 5.4 points
(D) 6.6 points
(E) 7.2 points

14. A multinational pharmaceutical company headquartered in Basel, Switzerland, is evaluating the effectiveness of a novel AI-powered diagnostic support tool designed to assist oncologists in treatment planning. The tool, branded "OncoAssist AI," analyzes patient medical records, genomic data, and clinical imaging to generate personalized treatment recommendations.
The company's Health Economics and Outcomes Research (HEOR) division conducted a twelve-month observational study across 45 cancer treatment centers. The primary outcome measure is the "Treatment Optimization Score" (TOS)—a composite metric ranging from 0 to 100.
Observed Adoption and Outcomes:

| Group | Mean TOS |
|---|---|
| OncoAssist AI Adopters (360) | 78.4 |
| Non-Adopters (540) | 62.8 |

Instrumental Variable: During the study period, 18 of the 45 treatment centers experienced a hospital-wide IT infrastructure upgrade mandated by regional health authorities. This upgrade made the installation of OncoAssist AI substantially easier. The timing was determined by regulatory schedules established three years prior—completely independent of oncologist characteristics.
Data Summary by IT Upgrade Status:

| | WITH IT Upgrade | WITHOUT IT Upgrade |
|---|---|---|
| Adoption rate | 54% | 28% |
| Mean TOS | 72.6 | 66.2 |

The LATE formula is: LATE = (Difference in outcomes by instrument) / (Difference in adoption rates by instrument)
Based on the instrumental variable analysis, what is the estimated true causal effect (LATE) of OncoAssist AI adoption on Treatment Optimization Score?
(A) 6.4 points
(B) 15.6 points
(C) 22.3 points
(D) 24.6 points
(E) 28.4 points

15. A direct-to-consumer e-commerce company headquartered in Los Angeles specializes in premium skincare products. The company's Chief Marketing Officer has developed a sophisticated framework for allocating marketing budgets across experimental advertising campaigns based on the principle of maximizing the expected logarithmic growth rate of customer acquisition capital.

接次頁

The marketing analytics team has identified an unconventional opportunity: a celebrity influencer with 15 million followers has offered an exclusive 24-hour promotional partnership. Based on historical data from similar influencer collaborations, the company's data scientists have analyzed engagement patterns, follower demographics, and conversion funnels to estimate the campaign's potential outcomes.

The analytics team presents their findings:

- Probability of campaign success: 75% (defined as achieving viral engagement that generates substantial new customer acquisition)
- Probability of campaign failure: 25% (defined as minimal engagement resulting in negligible customer acquisition)

The financial structure of the campaign operates as follows:

- If successful: The company's customer acquisition capital grows by 200% of the amount invested in the campaign (i.e., for every $1 invested, the company gains $2.00 in customer lifetime value, resulting in $3.00 total return).
- If unsuccessful: The company loses the entire amount invested in the campaign (i.e., the investment generates zero return and the capital is completely lost).

The company currently has a customer acquisition budget of $2 million available for experimental campaigns this quarter. Let f denote the fraction of this budget allocated to the influencer campaign, where $0 \leq f \leq 1$.

The CMO's objective is to choose f to maximize the expected logarithmic utility of terminal marketing capital:

$$\max \quad E[\ln(W_1)]$$

where $W_0$ represents the initial marketing budget and $W_1$ represents the marketing capital after the campaign outcome is realized.

The CMO explains to the marketing team: "If we invest fraction f of our budget and the campaign succeeds, our terminal capital becomes $W_0(1 + 2f)$—we keep our original capital plus gain 200% on the invested portion. If the campaign fails, our terminal capital becomes $W_0(1 - f)$—we lose the entire invested amount. Our goal is to find the optimal f* that maximizes our expected log-capital, which corresponds to maximizing our long-term compounded growth across many such campaign decisions."

A junior brand manager asks: "With a 75% success probability and 200% returns, shouldn't we invest aggressively—perhaps 75% or more of our budget?"

The CMO responds: "That intuition ignores the asymmetry between gains and losses in a compounding environment. The mathematics will reveal a more nuanced answer."

What is the optimal fraction f* of the marketing budget to allocate to this influencer campaign?

(A) 0.250
(B) 0.375
(C) 0.500
(D) 0.625
(E) 0.750

16. *The following is an excerpt from an internal memorandum circulated among the senior leadership team of Meridian Therapeutics, a mid-sized biopharmaceutical company headquartered in Cambridge, Massachusetts.*

In the fourth quarter of last fiscal year, Meridian Therapeutics initiated a pivotal Phase III clinical trial designated AURORA-7 to evaluate the efficacy of our lead compound, MRD-4821, for the treatment of treatment-resistant major depressive disorder (TR-MDD). The trial enrolled 840 patients across 62 clinical sites in North America and Western Europe, randomizing participants in a 1:1 ratio to receive either MRD-4821 or placebo over a 12-week treatment period. The primary endpoint was defined as the mean change from baseline in the Montgomery-Åsberg Depression Rating Scale (MADRS) score at Week 12.

The Data Safety Monitoring Board (DSMB) convened last week to review the unblinded interim analysis results. Dr. Helena Vasquez, the DSMB chair and a renowned biostatistician from Johns Hopkins University, presented the following summary to our executive committee:

見背面

"The interim analysis was conducted when approximately 60% of enrolled patients had completed the Week 12 assessment. At this juncture, patients receiving MRD-4821 demonstrated a mean MADRS reduction of 18.4 points compared to 14.2 points in the placebo arm, yielding a treatment difference of 4.2 points. The two-sided p-value associated with this comparison was 0.037. Per the pre-specified interim analysis plan, which employed an O'Brien-Fleming alpha-spending function, the boundary for declaring efficacy at this interim look was set at $p < 0.012$."

Following Dr. Vasquez's presentation, several members of our executive committee offered their interpretations of these results:

- Chief Executive Officer, Marcus Chen: "While I understand we haven't crossed the formal efficacy boundary, the p-value of 0.037 would be considered statistically significant under conventional standards. The DSMB's conservative threshold is merely a procedural technicality. I recommend we begin preparing our regulatory submission package immediately, as these results clearly demonstrate that MRD-4821 works."
- Chief Medical Officer, Dr. Priya Sharma: "I must respectfully disagree with Marcus's characterization. The alpha-spending approach exists precisely to preserve the overall Type I error rate across multiple analyses. If we had pre-specified only a single final analysis, a p-value of 0.037 would indeed cross the conventional 0.05 threshold. However, by conducting this interim look with the possibility of early stopping, we have effectively performed multiple hypothesis tests on accumulating data. The O'Brien-Fleming boundary of 0.012 at this interim analysis is not a 'technicality'—it is the mathematically correct threshold required to maintain a 5% overall false positive rate. Our current results, while encouraging, do not constitute definitive evidence of efficacy."
- Chief Financial Officer, Jonathan Mills: "From a capital markets perspective, I'm concerned about the investment community's reaction if these results were to become public prematurely. Our stock price has already appreciated 40% over the past six months in anticipation of positive AURORA-7 data. If we announce that the trial 'failed to meet its interim endpoint,' shareholders may interpret this as a negative signal, even though the trial is continuing to its final analysis. I suggest we refrain from any public disclosure until the final results are available."
- Chief Commercial Officer, Diana Rodriguez: "I've been analyzing the competitive landscape, and I believe timing is critical. Two of our competitors are expected to report Phase III results for their own TR-MDD candidates within the next nine months. If we wait for our final analysis—which won't be available for another eight months—we risk losing first-mover advantage in physician awareness and formulary positioning. The interim p-value of 0.037 suggests a very high probability that the final analysis will be positive. I recommend we accelerate our commercial launch preparations based on these interim findings."
- Vice President of Biostatistics, Dr. Samuel Okonkwo: "I want to address a point of confusion that seems to be emerging. Diana mentioned that the interim p-value 'suggests a very high probability that the final analysis will be positive.' This interpretation reflects a common misconception. A p-value of 0.037 means that, assuming the null hypothesis is true—that is, assuming MRD-4821 has no effect whatsoever—there would be a 3.7% probability of observing a treatment difference at least as large as what we found. It does not tell us the probability that MRD-4821 actually works, nor does it directly predict the outcome of the final analysis. The final analysis could yield a smaller treatment effect due to regression to the mean, differential dropout patterns, or site-by-treatment interactions that manifest differently as enrollment matures."

The DSMB has recommended that the trial continue to its pre-planned final analysis without modification. The executive committee must now decide on the company's strategic posture for the coming months. Based on the passage, which of the following statements represents the most accurate interpretation of the statistical evidence and reasoning presented?

接次頁

(A) The CEO's position is fundamentally sound because a p-value of 0.037 provides strong evidence that MRD-4821 is effective; the O'Brien-Fleming boundary is an arbitrary administrative threshold that does not affect the scientific interpretation of the data.

(B) The CFO's concern about investor reaction is the most statistically sophisticated perspective presented, as it recognizes that the market's interpretation of hypothesis testing outcomes may differ from the technical statistical conclusions.

(C) The CMO's reasoning is correct: the interim efficacy boundary of $p < 0.012$ is not merely procedural but is the appropriate threshold for this analysis, because conducting multiple analyses on accumulating data inflates the probability of a false positive finding unless stricter significance levels are applied at each look.

(D) The CCO's inference is valid because a p-value below 0.05 at the interim analysis mathematically guarantees at least a 95% probability that the final analysis will also yield a statistically significant result.

(E) Dr. Okonkwo's critique of the CCO's reasoning is flawed because p-values do, by definition, represent the probability that the experimental treatment is effective given the observed data.

17. Dr. Helena Morrison, a behavioral economist at the London School of Economics, is conducting a study on decision-making under uncertainty. She has designed an experimental maze for her research subjects. The maze has a single starting chamber with exactly three doors, labeled A, B, and C. Each door leads to a different tunnel, and the subjects have no information about where each tunnel leads. Due to the maze's construction, once a subject enters a tunnel, they cannot turn back and must complete that path.

The three tunnels have the following properties, unknown to the subjects: Door A leads through a winding passage that takes exactly 2 hours to traverse, after which the subject emerges at the exit and successfully completes the maze. Door B leads through a deceptive corridor that takes 3 hours to traverse, but ultimately returns the subject to the original starting chamber. Door C leads through a short loop that takes 1 hour to traverse, also returning the subject to the starting chamber.

Each time a subject arrives at (or returns to) the starting chamber, they choose among the three doors with equal probability of 1/3 each, as subjects cannot remember or distinguish which doors they have previously tried. Dr. Morrison wishes to calculate the expected time for a subject to exit the maze.

What is the expected number of hours for a subject to exit the maze?

(A) 4 hours

(B) 6 hours

(C) 8 hours

(D) 10 hours

18. Global Entertainment Network (GEN) is reviving a classic television game show format for the 2025 season. The show's signature segment works as follows: A contestant is presented with three identical doors on stage. Behind one door is a brand-new Tesla Model S valued at $89,000, while each of the other two doors conceals a decorative houseplant worth approximately $15. The contestant has no prior information about which door hides the car.

The game proceeds in two stages. In Stage 1, the contestant selects one of the three doors but does not open it. In Stage 2, the host—who knows exactly what is behind each door—deliberately opens one of the two remaining doors to reveal a houseplant. Crucially, the host never opens the door concealing the car, and if the contestant's initial choice was a houseplant, the host has only one valid door to open; if the contestant initially chose the car, the host randomly selects one of the two houseplant doors to open. After the host reveals a houseplant, the contestant is offered a choice: stay with the original selection or switch to the other unopened door.

The show's statistical consultant has been asked to determine the probability of winning the car under each strategy, assuming the contestant's initial choice is random.

What is the probability of winning the car if the contestant always switches, compared to always staying?

(A) Switching: 1/2; Staying: 1/2 (no advantage either way)

(B) Switching: 2/3; Staying: 1/3 (switching doubles the odds)

(C) Switching: 1/3; Staying: 2/3 (staying is better)

見背面

(D) Switching: 3/4; Staying: 1/4 (switching triples the odds)

19. The Taipei Metro operates the Wenhu Line with trains departing from Taipei Zoo station according to a Poisson process. On average, trains depart at a rate of $\lambda = 6$ trains per hour during peak hours, meaning the inter-arrival times between successive trains are independent and exponentially distributed with mean 10 minutes.

A commuter arrives at the platform at a uniformly random time during the peak hour (i.e., the commuter's arrival time is independent of the train schedule). Intuitively, one might expect the average waiting time for the next train to be half the average inter-arrival time, i.e., 5 minutes. However, a transportation analyst claims that the expected waiting time is actually longer due to a phenomenon known as the "inspection paradox" or "waiting time paradox."

The paradox arises because a randomly arriving passenger is more likely to arrive during a longer-than-average inter-arrival interval (since longer intervals occupy more time and thus have a higher probability of being "hit" by a random arrival). For a Poisson process, however, the memoryless property of the exponential distribution implies that the waiting time from any arbitrary arrival point has the same distribution as the original inter-arrival time.

Given that trains arrive according to a Poisson process with rate 6 per hour, what is the expected waiting time for a randomly arriving passenger?

(A) 5 minutes
(B) 10 minutes
(C) 15 minutes
(D) 20 minutes

20. A major technology company has implemented an AI-powered security system to detect unauthorized access attempts on its corporate network. The system was developed by analyzing historical data and has the following validated performance characteristics: When a genuine intrusion attempt occurs, the system correctly identifies it as an intrusion with probability 0.99 (true positive rate / sensitivity). When normal, authorized activity occurs, the system incorrectly flags it as an intrusion with probability 0.02 (false positive rate).

The company's cybersecurity team estimates that, on any given day, approximately 0.1% of all network activities represent genuine intrusion attempts, while 99.9% are normal authorized operations. On a particular Monday morning, the AI system flags a specific activity as a potential intrusion, triggering an alert to the security operations center.

The Chief Information Security Officer (CISO) asks the analytics team to calculate the probability that this flagged activity is actually a genuine intrusion, given that the system has raised an alert. This requires applying Bayes' theorem to update the prior probability of intrusion (0.1%) based on the evidence of a positive alert.

What is the probability that the flagged activity is a genuine intrusion, given that the AI system raised an alert?

(A) Approximately 1%
(B) Approximately 5%
(C) Approximately 33%
(D) Approximately 99%

接次頁

第二大題 (20%)

※ Show the detailed calculation process for all questions.

21. (10%) The length of time between calls received by the switchboard in a large legal firm is an exponential random variable. The average length of time between calls is 20 seconds. If a call has just been received, what is the probability that no calls are received in the next 30 seconds?

22. (10%) Suppose a surfboard company has found its yearly profit equation to be $P(x,y) = -2x^2 + 2xy - y^2 + 10x - 4y + 107$ , where $x$ is the number (in thousands) of type-A surfboards produced per year, $y$ is the number (in thousands) of type-B surfboards produced per year, and $P$ is profit (in thousands of dollars). How many of each type of board should be produced per year to realize a maximum profit? What is the maximum profit?

第三大題 (40%)：請依題號順序作答，可以用中文作答。

**Question I (20 points; 10 points each)**
In this question, we introduce autocorrelation in the disturbances but retain all the other assumptions of the classical model. The term autocorrelation may be defined as "correlation between members of series of observations ordered in time as in time series data." To explain the basic ideas involved, we consider the two-variable regression model

$$y_t = \beta_0 + \beta_1 x_t + u_t.$$

Note that we are using the subscript $t$ to emphasize that we are dealing with time series data. Assume that the disturbance terms follow the AR(1) scheme, which is known as a first-order autoregressive scheme:

$$u_t = \rho u_{t-1} + \varepsilon_t, \ -1 < \rho < 1,$$

where $\rho \neq 0$ is the coefficient of autocovariance and where $\varepsilon_t$ is the stochastic disturbance term such that it satisfies the standard ordinary least squares (OLS) assumptions, namely,

$$\begin{gathered} \mathrm{E}(\varepsilon_t) = 0 \\ \mathrm{var}(\varepsilon_t) = \sigma_\varepsilon^2 \\ \mathrm{cov}(\varepsilon_t, \varepsilon_{t+s}) = 0, \ s \neq 0. \end{gathered}$$

It can be shown that the properties of the disturbance term $u_t$ are

$$\mathrm{E}(u_t) = 0 \ \text{ and } \ \mathrm{var}(u_t) = \sigma_\varepsilon^2/(1-\rho^2).$$

23. Show that, under the AR(1) scheme, the disturbance term exhibits autocorrelation by deriving $\mathrm{E}(u_t u_{t-1}) \neq 0$.
24. Suppose that $\rho$ is known. Show that the original model can be transformed into the following form:

$$y_t = \alpha_0 + \alpha_1 x_t^* + \alpha_2 y_{t-1} + \varepsilon_t,$$

where $x_t^* = x_t - \rho x_{t-1}$. In particular, find the $\alpha_i$ in terms of $\beta_i$ and $\rho$.

見背面

**Question II (20 points; 5 points each)**

In a study prepared for the U.S. Department of Labor, researchers were interested in examining the labor-force participation of various categories of labor as a function of several explanatory variables. In all their regressions, the dependent variable is a dummy, taking a value of 1 if a person is in the labor force, 0 if she is not. Table II reports one of the estimated regressions, which is obtained using ordinary least squares (OLS).

Turning to interpretations of the findings, we see that each slope coefficient gives the rate of change in the conditional probability of the event occurring for a given unit change in the value of the explanatory variable. For instance, the coefficient of -0.27 attached to the variable "age 65 and over" means, holding all other factors constant, the probability of participation in the labor force by women in this age group is smaller by about 27 percent (as compared with the base category of women aged 22 to 54). Now consider the interaction term marital status and age. Table II shows that the probability of participation of women who were never married and are 65 or over is smaller by about 20 percent as compared with the base category. Following this procedure, readers can easily interpret the rest of the coefficients given in Table II.

Table II. Labor-Force Participation
Regression of women, age 22 and over, living in largest 96 standard metropolitan statistical areas (SMSA) (dependent variable: in or out of labor force during 1966)

| Explanatory Variable | | Coefficient | *t* Ratio |
|---|---|---|---|
| Constant | | 0.24 | 15.4 |
| Marital status (dummy) | | | |
| Married | | — | — |
| Never married | | 0.29 | 22.0 |
| Age (dummy) | | | |
| 22-54 | | — | — |
| 55-64 | | -0.06 | -5.7 |
| 65 and over | | -0.27 | -9.0 |
| Years of schooling | | 0.02 | 5.8 |
| Interaction (marital status and age) | | | |
| Marital status | Age | | |
| Never married | 55-64 | -0.11 | -3.3 |
| Never married | 65 and over | -0.20 | -6.4 |

$R^2 = 0.175$
No. of observations = 25,153
Note: — indicates the base or omitted category.

25. Based on the regression results, interpret the effect of education on women's labor force participation and determine whether the effect is statistically significant.
26. One would like to obtain the estimates of the conditional probabilities of labor-force participation of the various categories. In particular, compute the probability for women who were never married, aged 55 to 64, with 16 years of schooling.
27. By how many percentage points is the labor-force participation probability higher or lower for women who were never married and aged 55-64 compared with those who were never married and are 65 or over?
28. How would you test the hypothesis that the effect of married status on labor-force participation probability does not depend on age? State the hypotheses and the appropriate test.

**試題隨卷繳回**