The Economics of Proof-of-Stake Security – Quiz

20 Multiple-Choice Questions

Bloom's levels: 4 Understand · 8 Apply · 6 Analyze · 2 Evaluate

Prof. Dr. Jörg Osterrieder

BSc Blockchain, Crypto Economy & NFTs

Spring 2026

Quiz Questions 1–5

Q1. What distinguishes "cost-of-attack" from "cost-of-corruption" in blockchain security?
A) Cost-of-attack measures hardware; cost-of-corruption measures energy B) Cost-of-attack is execution cost; cost-of-corruption is the capital permanently destroyed on detection
C) They are equivalent terms for the same concept D) Cost-of-attack applies to PoS; cost-of-corruption applies to PoW

Quiz Questions 1–5

Q1. What distinguishes "cost-of-attack" from "cost-of-corruption" in blockchain security?
A) Cost-of-attack measures hardware; cost-of-corruption measures energy B) Cost-of-attack is execution cost; cost-of-corruption is the capital permanently destroyed on detection
C) They are equivalent terms for the same concept D) Cost-of-attack applies to PoS; cost-of-corruption applies to PoW
Answer: B – Zamfir (2015) defined cost-of-corruption as the provable capital destroyed when an attack is detected, distinct from the upfront execution cost.

Q2. An attacker holds 40% of staked ETH. Using $S = 3f \times$ Bond, what fraction of the bond is burned on detection?
A) 40% of the bond, proportional to the attacker's share of total staked ETH B) 0% (penalty only applies to individual validators) C) 120% of their initial bond (capped at 100%) D) 33% of their bond at a fixed rate set by the Ethereum protocol

Quiz Questions 1–5

Q1. What distinguishes "cost-of-attack" from "cost-of-corruption" in blockchain security?
A) Cost-of-attack measures hardware; cost-of-corruption measures energy B) Cost-of-attack is execution cost; cost-of-corruption is the capital permanently destroyed on detection
C) They are equivalent terms for the same concept D) Cost-of-attack applies to PoS; cost-of-corruption applies to PoW
Answer: B – Zamfir (2015) defined cost-of-corruption as the provable capital destroyed when an attack is detected, distinct from the upfront execution cost.

Q2. An attacker holds 40% of staked ETH. Using $S = 3f \times$ Bond, what fraction of the bond is burned on detection?
A) 40% of the bond, proportional to the attacker's share of total staked ETH B) 0% (penalty only applies to individual validators) C) 120% of their initial bond (capped at 100%) D) 33% of their bond at a fixed rate set by the Ethereum protocol
Answer: C – $3 \times 0.40 = 1.20$: the correlation penalty exceeds the bond, so 100% is burned (capped at the bond value).

Q3. Why is PoW attack capital "rentable" while PoS attack capital is "non-rentable"?
A) PoW miners pay rent for block space; PoS validators own their stake B) PoW hardware is owned by mining pools; PoS validators own equipment individually
C) PoS validators can rent out stake for income; PoW miners cannot D) PoW hash power can be leased via NiceHash without owning hardware; PoS attack requires buying and bonding ETH outright

Quiz Questions 1–5

Q1. What distinguishes "cost-of-attack" from "cost-of-corruption" in blockchain security?
A) Cost-of-attack measures hardware; cost-of-corruption measures energy B) Cost-of-attack is execution cost; cost-of-corruption is the capital permanently destroyed on detection
C) They are equivalent terms for the same concept D) Cost-of-attack applies to PoS; cost-of-corruption applies to PoW
Answer: B – Zamfir (2015) defined cost-of-corruption as the provable capital destroyed when an attack is detected, distinct from the upfront execution cost.

Q2. An attacker holds 40% of staked ETH. Using $S = 3f \times$ Bond, what fraction of the bond is burned on detection?
A) 40% of the bond, proportional to the attacker's share of total staked ETH B) 0% (penalty only applies to individual validators) C) 120% of their initial bond (capped at 100%) D) 33% of their bond at a fixed rate set by the Ethereum protocol
Answer: C – $3 \times 0.40 = 1.20$: the correlation penalty exceeds the bond, so 100% is burned (capped at the bond value).

Q3. Why is PoW attack capital "rentable" while PoS attack capital is "non-rentable"?
A) PoW miners pay rent for block space; PoS validators own their stake B) PoW hardware is owned by mining pools; PoS validators own equipment individually
C) PoS validators can rent out stake for income; PoW miners cannot D) PoW hash power can be leased via NiceHash without owning hardware; PoS attack requires buying and bonding ETH outright
Answer: D – NiceHash and similar platforms allow temporary hash power rental with no permanent commitment. No equivalent attack-capital lease market exists for PoS.

Q4. A PoW attacker spends $10M on ASICs and fails to execute a 51% attack. Approximately what fraction of the $10M can be recovered?
A) Approximately 60-70% via the hardware resale market B) 0% (hardware worthless after failure) C) Exactly 33% (standard ASIC depreciation rate) D) 100% (ASICs retain full value)

Quiz Questions 1–5

Q1. What distinguishes "cost-of-attack" from "cost-of-corruption" in blockchain security?
A) Cost-of-attack measures hardware; cost-of-corruption measures energy B) Cost-of-attack is execution cost; cost-of-corruption is the capital permanently destroyed on detection
C) They are equivalent terms for the same concept D) Cost-of-attack applies to PoS; cost-of-corruption applies to PoW
Answer: B – Zamfir (2015) defined cost-of-corruption as the provable capital destroyed when an attack is detected, distinct from the upfront execution cost.

Q2. An attacker holds 40% of staked ETH. Using $S = 3f \times$ Bond, what fraction of the bond is burned on detection?
A) 40% of the bond, proportional to the attacker's share of total staked ETH B) 0% (penalty only applies to individual validators) C) 120% of their initial bond (capped at 100%) D) 33% of their bond at a fixed rate set by the Ethereum protocol
Answer: C – $3 \times 0.40 = 1.20$: the correlation penalty exceeds the bond, so 100% is burned (capped at the bond value).

Q3. Why is PoW attack capital "rentable" while PoS attack capital is "non-rentable"?
A) PoW miners pay rent for block space; PoS validators own their stake B) PoW hardware is owned by mining pools; PoS validators own equipment individually
C) PoS validators can rent out stake for income; PoW miners cannot D) PoW hash power can be leased via NiceHash without owning hardware; PoS attack requires buying and bonding ETH outright
Answer: D – NiceHash and similar platforms allow temporary hash power rental with no permanent commitment. No equivalent attack-capital lease market exists for PoS.

Q4. A PoW attacker spends $10M on ASICs and fails to execute a 51% attack. Approximately what fraction of the $10M can be recovered?
A) Approximately 60-70% via the hardware resale market B) 0% (hardware worthless after failure) C) Exactly 33% (standard ASIC depreciation rate) D) 100% (ASICs retain full value)
Answer: A – ASICs retain substantial market value and can be resold, unlike a burned PoS bond.

Q5. In the expected payoff formula $\mathbb{E}[\text{payoff}] = p \cdot G - (1-p) \cdot S - C_b$, what does S represent?
A) The probability that the attack succeeds on the first attempt B) The financial gain from a successful double-spend or reorg attack C) The slashed bond burned by the protocol on detection D) The staking reward earned by fully honest non-attacking validators per epoch

Quiz Questions 1–5

Q1. What distinguishes "cost-of-attack" from "cost-of-corruption" in blockchain security?
A) Cost-of-attack measures hardware; cost-of-corruption measures energy B) Cost-of-attack is execution cost; cost-of-corruption is the capital permanently destroyed on detection
C) They are equivalent terms for the same concept D) Cost-of-attack applies to PoS; cost-of-corruption applies to PoW
Answer: B – Zamfir (2015) defined cost-of-corruption as the provable capital destroyed when an attack is detected, distinct from the upfront execution cost.

Q2. An attacker holds 40% of staked ETH. Using $S = 3f \times$ Bond, what fraction of the bond is burned on detection?
A) 40% of the bond, proportional to the attacker's share of total staked ETH B) 0% (penalty only applies to individual validators) C) 120% of their initial bond (capped at 100%) D) 33% of their bond at a fixed rate set by the Ethereum protocol
Answer: C – $3 \times 0.40 = 1.20$: the correlation penalty exceeds the bond, so 100% is burned (capped at the bond value).

Q3. Why is PoW attack capital "rentable" while PoS attack capital is "non-rentable"?
A) PoW miners pay rent for block space; PoS validators own their stake B) PoW hardware is owned by mining pools; PoS validators own equipment individually
C) PoS validators can rent out stake for income; PoW miners cannot D) PoW hash power can be leased via NiceHash without owning hardware; PoS attack requires buying and bonding ETH outright
Answer: D – NiceHash and similar platforms allow temporary hash power rental with no permanent commitment. No equivalent attack-capital lease market exists for PoS.

Q4. A PoW attacker spends \$10M on ASICs and fails to execute a 51% attack. Approximately what fraction of the \$10M can be recovered?
A) Approximately 60-70% via the hardware resale market B) 0% (hardware worthless after failure) C) Exactly 33% (standard ASIC depreciation rate) D) 100% (ASICs retain full value)
Answer: A – ASICs retain substantial market value and can be resold, unlike a burned PoS bond.

Q5. In the expected payoff formula $\mathbb{E}[\text{payoff}] = p \cdot G - (1-p) \cdot S - C_b$, what does S represent?
A) The probability that the attack succeeds on the first attempt B) The financial gain from a successful double-spend or reorg attack C) The slashed bond burned by the protocol on detection D) The staking reward earned by fully honest non-attacking validators per epoch
Answer: C – S is the slashed bond destroyed by the correlation penalty on detection; it is the dominant term at high stake fractions.

Quiz Questions 6–10

Q6. Using $S = 3f \times$ Bond, what fraction of the bond is burned for an attacker holding exactly 10% of total stake?
A) 10% of their bond B) 100% of their bond C) 133% of their bond D) 30% of their bond

Quiz Questions 6–10

Q6. Using $S = 3f \times$ Bond, what fraction of the bond is burned for an attacker holding exactly 10% of total stake?
A) 10% of their bond B) 100% of their bond C) 133% of their bond D) 30% of their bond
Answer: D – $3 \times 0.10 = 0.30$, so 30% of the bond is burned. The correlation penalty is non-linear: it grows faster than the attacker's stake.

Q7. For which attacker stake fraction f does the correlation penalty S equal exactly the full bond?
A) $f = 10\%$ B) $f = 20\%$ C) $f = 33\%$ D) $f = 50\%$

Quiz Questions 6–10

Q6. Using $S = 3f \times$ Bond, what fraction of the bond is burned for an attacker holding exactly 10% of total stake?
A) 10% of their bond B) 100% of their bond C) 133% of their bond D) 30% of their bond
Answer: D – $3 \times 0.10 = 0.30$, so 30% of the bond is burned. The correlation penalty is non-linear: it grows faster than the attacker's stake.

Q7. For which attacker stake fraction f does the correlation penalty S equal exactly the full bond?
A) $f = 10\%$ B) $f = 20\%$ C) $f = 33\%$ D) $f = 50\%$
Answer: C – $3 \times 0.33 \approx 1.00$: at 33% the penalty equals 100% of the bond, meaning full capital loss.

Q8. Why does the expected payoff of a PoS attack grow more negative as the attacker accumulates more stake?
A) The slashed amount $S = 3f \times$ Bond grows non-linearly while gain G is bounded, so S outpaces G at higher stake levels B) Larger attackers earn smaller staking rewards
C) Detection probability p falls as the attacker accumulates a larger share of the validator committee, because a coordinated majority can suppress slashing evidence by refusing to include it in blocks D) The borrow cost C_b falls with larger stake

Quiz Questions 6–10

Q6. Using $S = 3f \times$ Bond, what fraction of the bond is burned for an attacker holding exactly 10% of total stake?
A) 10% of their bond B) 100% of their bond C) 133% of their bond D) 30% of their bond
Answer: D – $3 \times 0.10 = 0.30$, so 30% of the bond is burned. The correlation penalty is non-linear: it grows faster than the attacker's stake.

Q7. For which attacker stake fraction f does the correlation penalty S equal exactly the full bond?
A) $f = 10\%$ B) $f = 20\%$ C) $f = 33\%$ D) $f = 50\%$
Answer: C – $3 \times 0.33 \approx 1.00$: at 33% the penalty equals 100% of the bond, meaning full capital loss.

Q8. Why does the expected payoff of a PoS attack grow more negative as the attacker accumulates more stake?
A) The slashed amount $S = 3f \times$ Bond grows non-linearly while gain G is bounded, so S outpaces G at higher stake levels B) Larger attackers earn smaller staking rewards
C) Detection probability p falls as the attacker accumulates a larger share of the validator committee, because a coordinated majority can suppress slashing evidence by refusing to include it in blocks D) The borrow cost C_b falls with larger stake
Answer: A – S grows with $3f$, while G (double-spend profit) is bounded by market liquidity. The tighter the inequality, the more irrational the attack.

Q9. What does "economic finality" mean in PoS, versus "probabilistic finality" in PoW?
A) PoS reversion costs $\geq$33% of staked ETH; PoW reversal risk falls exponentially but never reaches exactly zero B) PoS confirms faster in all cases C) PoS finality is decided by economic votes; PoW finality is random D) PoS and PoW both provide deterministic finality; the difference is only how block rewards are distributed to participants

Quiz Questions 6–10

Q6. Using $S = 3f \times$ Bond, what fraction of the bond is burned for an attacker holding exactly 10% of total stake?
A) 10% of their bond B) 100% of their bond C) 133% of their bond D) 30% of their bond
Answer: D – $3 \times 0.10 = 0.30$, so 30% of the bond is burned. The correlation penalty is non-linear: it grows faster than the attacker's stake.

Q7. For which attacker stake fraction f does the correlation penalty S equal exactly the full bond?
A) $f = 10\%$ B) $f = 20\%$ C) $f = 33\%$ D) $f = 50\%$
Answer: C – $3 \times 0.33 \approx 1.00$: at 33% the penalty equals 100% of the bond, meaning full capital loss.

Q8. Why does the expected payoff of a PoS attack grow more negative as the attacker accumulates more stake?
A) The slashed amount $S = 3f \times$ Bond grows non-linearly while gain G is bounded, so S outpaces G at higher stake levels B) Larger attackers earn smaller staking rewards
C) Detection probability p falls as the attacker accumulates a larger share of the validator committee, because a coordinated majority can suppress slashing evidence by refusing to include it in blocks D) The borrow cost C_b falls with larger stake
Answer: A – S grows with $3f$, while G (double-spend profit) is bounded by market liquidity. The tighter the inequality, the more irrational the attack.

Q9. What does "economic finality" mean in PoS, versus "probabilistic finality" in PoW?
A) PoS reversion costs $\geq$33% of staked ETH; PoW reversal risk falls exponentially but never reaches exactly zero B) PoS confirms faster in all cases C) PoS finality is decided by economic votes; PoW finality is random D) PoS and PoW both provide deterministic finality; the difference is only how block rewards are distributed to participants
Answer: A – Casper FFG makes reversion provably costly (the attacker loses 33%+ of staked ETH). Nakamoto finality is probabilistic: deep reversal is expensive but never impossible in theory.

Q10. If $p = 0.95$ and $G = \$10$M, what must S exceed for the attack to be irrational (ignoring C_b)?
A) $S > \$10$M B) $S > \$1$M C) $S > \$190$M D) $S > \$1{,}900$M

Quiz Questions 6–10

Q6. Using $S = 3f \times$ Bond, what fraction of the bond is burned for an attacker holding exactly 10% of total stake?
A) 10% of their bond B) 100% of their bond C) 133% of their bond D) 30% of their bond
Answer: D – $3 \times 0.10 = 0.30$, so 30% of the bond is burned. The correlation penalty is non-linear: it grows faster than the attacker's stake.

Q7. For which attacker stake fraction f does the correlation penalty S equal exactly the full bond?
A) $f = 10\%$ B) $f = 20\%$ C) $f = 33\%$ D) $f = 50\%$
Answer: C – $3 \times 0.33 \approx 1.00$: at 33% the penalty equals 100% of the bond, meaning full capital loss.

Q8. Why does the expected payoff of a PoS attack grow more negative as the attacker accumulates more stake?
A) The slashed amount $S = 3f \times$ Bond grows non-linearly while gain G is bounded, so S outpaces G at higher stake levels B) Larger attackers earn smaller staking rewards
C) Detection probability p falls as the attacker accumulates a larger share of the validator committee, because a coordinated majority can suppress slashing evidence by refusing to include it in blocks D) The borrow cost C_b falls with larger stake
Answer: A – S grows with $3f$, while G (double-spend profit) is bounded by market liquidity. The tighter the inequality, the more irrational the attack.

Q9. What does "economic finality" mean in PoS, versus "probabilistic finality" in PoW?
A) PoS reversion costs $\geq$33% of staked ETH; PoW reversal risk falls exponentially but never reaches exactly zero B) PoS confirms faster in all cases C) PoS finality is decided by economic votes; PoW finality is random D) PoS and PoW both provide deterministic finality; the difference is only how block rewards are distributed to participants
Answer: A – Casper FFG makes reversion provably costly (the attacker loses 33%+ of staked ETH). Nakamoto finality is probabilistic: deep reversal is expensive but never impossible in theory.

Q10. If $p = 0.95$ and $G =$ \$10M, what must S exceed for the attack to be irrational (ignoring C_b)?
A) $S >$ \$10M B) $S >$ \$1M C) $S >$ \$190M D) $S >$ \$1,900M
Answer: C – For $\mathbb{E}[\text{payoff}] < 0$: $0.95 \times 10 < 0.05 \times S$, giving $S >$ \$190M. This magnitude is achievable for large validators at 33%+ stake.

Quiz Questions 11–15

Q11. What is "weak subjectivity" in Proof-of-Stake?
A) Validators can subjectively choose which transactions to include B) PoS networks assign higher voting weight to validators who have a longer continuous record of supporting the majority chain, creating subjective advantage for long-standing validators over newcomers
C) Validators must subjectively confirm blocks before finalization D) A syncing node must get a checkpoint because it cannot distinguish the honest chain from a long-range fork by protocol rules alone

Quiz Questions 11–15

Q11. What is "weak subjectivity" in Proof-of-Stake?
A) Validators can subjectively choose which transactions to include B) PoS networks assign higher voting weight to validators who have a longer continuous record of supporting the majority chain, creating subjective advantage for long-standing validators over newcomers
C) Validators must subjectively confirm blocks before finalization D) A syncing node must get a checkpoint because it cannot distinguish the honest chain from a long-range fork by protocol rules alone
Answer: D – Named by Buterin (2014): a new client needs one piece of external social information (the checkpoint) beyond what the protocol can prove from genesis.

Q12. Why does PoW avoid the weak subjectivity problem?
A) PoW nodes use checkpoint data from trusted mining pool operators to verify the chain they download matches the current canonical head B) PoW requires real energy; a fake chain cannot match cumulative work without equivalent electricity expenditure C) PoW validators are randomly selected D) PoW chains are shorter than PoS chains

Quiz Questions 11–15

Q11. What is "weak subjectivity" in Proof-of-Stake?
A) Validators can subjectively choose which transactions to include B) PoS networks assign higher voting weight to validators who have a longer continuous record of supporting the majority chain, creating subjective advantage for long-standing validators over newcomers
C) Validators must subjectively confirm blocks before finalization D) A syncing node must get a checkpoint because it cannot distinguish the honest chain from a long-range fork by protocol rules alone
Answer: D – Named by Buterin (2014): a new client needs one piece of external social information (the checkpoint) beyond what the protocol can prove from genesis.

Q12. Why does PoW avoid the weak subjectivity problem?
A) PoW nodes use checkpoint data from trusted mining pool operators to verify the chain they download matches the current canonical head B) PoW requires real energy; a fake chain cannot match cumulative work without equivalent electricity expenditure C) PoW validators are randomly selected D) PoW chains are shorter than PoS chains
Answer: B – Energy expenditure makes PoW chain history unforgeable. In PoS, past block production was costless: old keys can simulate alternate histories without any energy cost.

Q13. An attacker uses keys from a validator set that withdrew 3 months ago. What attack can they mount against a newly syncing PoS node?
A) A 51% attack using current mining power B) A nothing-at-stake attack on current epoch validators C) A long-range attack using old keys to build an alternate chain history from 3 months ago that appears equally valid to a naive syncing node D) A slashing attack that reports currently bonded validators for double-signing, triggering the correlation penalty and burning their entire bond while leaving the attacker's own stake intact

Quiz Questions 11–15

Q11. What is "weak subjectivity" in Proof-of-Stake?
A) Validators can subjectively choose which transactions to include B) PoS networks assign higher voting weight to validators who have a longer continuous record of supporting the majority chain, creating subjective advantage for long-standing validators over newcomers
C) Validators must subjectively confirm blocks before finalization D) A syncing node must get a checkpoint because it cannot distinguish the honest chain from a long-range fork by protocol rules alone
Answer: D – Named by Buterin (2014): a new client needs one piece of external social information (the checkpoint) beyond what the protocol can prove from genesis.

Q12. Why does PoW avoid the weak subjectivity problem?
A) PoW nodes use checkpoint data from trusted mining pool operators to verify the chain they download matches the current canonical head B) PoW requires real energy; a fake chain cannot match cumulative work without equivalent electricity expenditure C) PoW validators are randomly selected D) PoW chains are shorter than PoS chains
Answer: B – Energy expenditure makes PoW chain history unforgeable. In PoS, past block production was costless: old keys can simulate alternate histories without any energy cost.

Q13. An attacker uses keys from a validator set that withdrew 3 months ago. What attack can they mount against a newly syncing PoS node?
A) A 51% attack using current mining power B) A nothing-at-stake attack on current epoch validators C) A long-range attack using old keys to build an alternate chain history from 3 months ago that appears equally valid to a naive syncing node D) A slashing attack that reports currently bonded validators for double-signing, triggering the correlation penalty and burning their entire bond while leaving the attacker's own stake intact
Answer: C – Withdrawn validators retain their old signing keys. Checkpoints defeat this by providing an anchor more recent than the fork point.

Q14. What is the approximate weak subjectivity period on Ethereum mainnet (as of 2023)?
A) 10 minutes (one epoch) B) 6 hours (one finality slot) C) Approximately 2 weeks D) 1 year (withdrawal delay)

Quiz Questions 11–15

Q11. What is "weak subjectivity" in Proof-of-Stake?
A) Validators can subjectively choose which transactions to include B) PoS networks assign higher voting weight to validators who have a longer continuous record of supporting the majority chain, creating subjective advantage for long-standing validators over newcomers
C) Validators must subjectively confirm blocks before finalization D) A syncing node must get a checkpoint because it cannot distinguish the honest chain from a long-range fork by protocol rules alone
Answer: D – Named by Buterin (2014): a new client needs one piece of external social information (the checkpoint) beyond what the protocol can prove from genesis.

Q12. Why does PoW avoid the weak subjectivity problem?
A) PoW nodes use checkpoint data from trusted mining pool operators to verify the chain they download matches the current canonical head B) PoW requires real energy; a fake chain cannot match cumulative work without equivalent electricity expenditure C) PoW validators are randomly selected D) PoW chains are shorter than PoS chains
Answer: B – Energy expenditure makes PoW chain history unforgeable. In PoS, past block production was costless: old keys can simulate alternate histories without any energy cost.

Q13. An attacker uses keys from a validator set that withdrew 3 months ago. What attack can they mount against a newly syncing PoS node?
A) A 51% attack using current mining power B) A nothing-at-stake attack on current epoch validators C) A long-range attack using old keys to build an alternate chain history from 3 months ago that appears equally valid to a naive syncing node D) A slashing attack that reports currently bonded validators for double-signing, triggering the correlation penalty and burning their entire bond while leaving the attacker's own stake intact
Answer: C – Withdrawn validators retain their old signing keys. Checkpoints defeat this by providing an anchor more recent than the fork point.

Q14. What is the approximate weak subjectivity period on Ethereum mainnet (as of 2023)?
A) 10 minutes (one epoch) B) 6 hours (one finality slot) C) Approximately 2 weeks D) 1 year (withdrawal delay)
Answer: C – The Ethereum Foundation documentation specifies approximately 2 weeks; nodes offline longer than this must obtain a recent checkpoint before syncing.

Q15. A user running an Ethereum node has been offline for 3 weeks. What must they do before trusting the chain?
A) Obtain a recent finalized checkpoint from a trusted source (Ethereum Foundation portal, known peers) before resuming sync B) Nothing; the longest chain is automatically valid C) Wait for the next protocol upgrade to reset finality D) Re-download the full blockchain from genesis and replay all transactions from scratch, since the protocol can always identify the canonical chain by reprocessing history without any external checkpoint

Quiz Questions 11–15

Q11. What is "weak subjectivity" in Proof-of-Stake?
A) Validators can subjectively choose which transactions to include B) PoS networks assign higher voting weight to validators who have a longer continuous record of supporting the majority chain, creating subjective advantage for long-standing validators over newcomers
C) Validators must subjectively confirm blocks before finalization D) A syncing node must get a checkpoint because it cannot distinguish the honest chain from a long-range fork by protocol rules alone
Answer: D – Named by Buterin (2014): a new client needs one piece of external social information (the checkpoint) beyond what the protocol can prove from genesis.

Q12. Why does PoW avoid the weak subjectivity problem?
A) PoW nodes use checkpoint data from trusted mining pool operators to verify the chain they download matches the current canonical head B) PoW requires real energy; a fake chain cannot match cumulative work without equivalent electricity expenditure C) PoW validators are randomly selected D) PoW chains are shorter than PoS chains
Answer: B – Energy expenditure makes PoW chain history unforgeable. In PoS, past block production was costless: old keys can simulate alternate histories without any energy cost.

Q13. An attacker uses keys from a validator set that withdrew 3 months ago. What attack can they mount against a newly syncing PoS node?
A) A 51% attack using current mining power B) A nothing-at-stake attack on current epoch validators C) A long-range attack using old keys to build an alternate chain history from 3 months ago that appears equally valid to a naive syncing node D) A slashing attack that reports currently bonded validators for double-signing, triggering the correlation penalty and burning their entire bond while leaving the attacker's own stake intact
Answer: C – Withdrawn validators retain their old signing keys. Checkpoints defeat this by providing an anchor more recent than the fork point.

Q14. What is the approximate weak subjectivity period on Ethereum mainnet (as of 2023)?
A) 10 minutes (one epoch) B) 6 hours (one finality slot) C) Approximately 2 weeks D) 1 year (withdrawal delay)
Answer: C – The Ethereum Foundation documentation specifies approximately 2 weeks; nodes offline longer than this must obtain a recent checkpoint before syncing.

Q15. A user running an Ethereum node has been offline for 3 weeks. What must they do before trusting the chain?
A) Obtain a recent finalized checkpoint from a trusted source (Ethereum Foundation portal, known peers) before resuming sync B) Nothing; the longest chain is automatically valid C) Wait for the next protocol upgrade to reset finality D) Re-download the full blockchain from genesis and replay all transactions from scratch, since the protocol can always identify the canonical chain by reprocessing history without any external checkpoint
Answer: A – Three weeks exceeds the 2-week weak subjectivity period. The node must obtain a recent checkpoint to safely identify the honest chain.

Quiz Questions 16–20

Q16. Which dimension does this supplement add to L10's PoW/PoS consensus comparison?
A) Transaction throughput measured in confirmed transactions per second, since higher throughput implies more economic activity secured per block B) Energy consumption per transaction C) Block time and confirmation speed D) Whether attack capital can be recovered after a failed or detected attack

Quiz Questions 16–20

Q16. Which dimension does this supplement add to L10's PoW/PoS consensus comparison?
A) Transaction throughput measured in confirmed transactions per second, since higher throughput implies more economic activity secured per block B) Energy consumption per transaction C) Block time and confirmation speed D) Whether attack capital can be recovered after a failed or detected attack
Answer: D – L10 covers energy/TPS/finality speed. This deck adds: capital recoverability, finality type (probabilistic vs. economic), and attack-capital assembly time.

Q17. What is the primary ongoing cost of PoS network security, as opposed to PoW?
A) ETH issuance (inflation), paid collectively by all ETH holders as dilution B) Hardware depreciation, paid by validator operators C) Gas fees, paid by transaction senders D) Electricity costs billed to each validator operator for the continuous power consumption required to run hashing hardware around the clock and compete in the proof-of-work lottery

Quiz Questions 16–20

Q16. Which dimension does this supplement add to L10's PoW/PoS consensus comparison?
A) Transaction throughput measured in confirmed transactions per second, since higher throughput implies more economic activity secured per block B) Energy consumption per transaction C) Block time and confirmation speed D) Whether attack capital can be recovered after a failed or detected attack
Answer: D – L10 covers energy/TPS/finality speed. This deck adds: capital recoverability, finality type (probabilistic vs. economic), and attack-capital assembly time.

Q17. What is the primary ongoing cost of PoS network security, as opposed to PoW?
A) ETH issuance (inflation), paid collectively by all ETH holders as dilution B) Hardware depreciation, paid by validator operators C) Gas fees, paid by transaction senders D) Electricity costs billed to each validator operator for the continuous power consumption required to run hashing hardware around the clock and compete in the proof-of-work lottery
Answer: A – PoS validators earn newly issued ETH (inflation) as rewards. PoW miners earn block rewards plus fees but require continuous energy expenditure.

Q18. Why does a would-be PoS attacker face more entry friction than a PoW attacker?
A) PoS requires the private keys of the Ethereum Foundation B) PoW attacks require cooperation from all mining pools C) PoS validator selection uses verifiable randomness that prevents any party from predicting which validators will attest to specific blocks D) PoS attack capital must be bonded on-chain (slow, visible); PoW hash power can be rented via NiceHash in hours

Quiz Questions 16–20

Q16. Which dimension does this supplement add to L10's PoW/PoS consensus comparison?
A) Transaction throughput measured in confirmed transactions per second, since higher throughput implies more economic activity secured per block B) Energy consumption per transaction C) Block time and confirmation speed D) Whether attack capital can be recovered after a failed or detected attack
Answer: D – L10 covers energy/TPS/finality speed. This deck adds: capital recoverability, finality type (probabilistic vs. economic), and attack-capital assembly time.

Q17. What is the primary ongoing cost of PoS network security, as opposed to PoW?
A) ETH issuance (inflation), paid collectively by all ETH holders as dilution B) Hardware depreciation, paid by validator operators C) Gas fees, paid by transaction senders D) Electricity costs billed to each validator operator for the continuous power consumption required to run hashing hardware around the clock and compete in the proof-of-work lottery
Answer: A – PoS validators earn newly issued ETH (inflation) as rewards. PoW miners earn block rewards plus fees but require continuous energy expenditure.

Q18. Why does a would-be PoS attacker face more entry friction than a PoW attacker?
A) PoS requires the private keys of the Ethereum Foundation B) PoW attacks require cooperation from all mining pools C) PoS validator selection uses verifiable randomness that prevents any party from predicting which validators will attest to specific blocks D) PoS attack capital must be bonded on-chain (slow, visible); PoW hash power can be rented via NiceHash in hours
Answer: D – Assembling 33%+ of staked ETH is slow and publicly visible on-chain. NiceHash allows PoW attackers to assemble attack-scale hash power in hours with no on-chain trace.

Q19. A PoS attacker holds 35% of staked ETH. The attack is detected and the correlation penalty applied. What is their capital position?
A) They recover 65% of bonded ETH B) They lose exactly 35% of their bond C) They lose nothing (penalty only applies to double-signing within a single slot) D) They lose their entire bond: $3 \times 0.35 \times \text{Bond} = 1.05 \times \text{Bond}$, capped at 100%

Quiz Questions 16–20

Q16. Which dimension does this supplement add to L10's PoW/PoS consensus comparison?
A) Transaction throughput measured in confirmed transactions per second, since higher throughput implies more economic activity secured per block B) Energy consumption per transaction C) Block time and confirmation speed D) Whether attack capital can be recovered after a failed or detected attack
Answer: D – L10 covers energy/TPS/finality speed. This deck adds: capital recoverability, finality type (probabilistic vs. economic), and attack-capital assembly time.

Q17. What is the primary ongoing cost of PoS network security, as opposed to PoW?
A) ETH issuance (inflation), paid collectively by all ETH holders as dilution B) Hardware depreciation, paid by validator operators C) Gas fees, paid by transaction senders D) Electricity costs billed to each validator operator for the continuous power consumption required to run hashing hardware around the clock and compete in the proof-of-work lottery
Answer: A – PoS validators earn newly issued ETH (inflation) as rewards. PoW miners earn block rewards plus fees but require continuous energy expenditure.

Q18. Why does a would-be PoS attacker face more entry friction than a PoW attacker?
A) PoS requires the private keys of the Ethereum Foundation B) PoW attacks require cooperation from all mining pools C) PoS validator selection uses verifiable randomness that prevents any party from predicting which validators will attest to specific blocks D) PoS attack capital must be bonded on-chain (slow, visible); PoW hash power can be rented via NiceHash in hours
Answer: D – Assembling 33%+ of staked ETH is slow and publicly visible on-chain. NiceHash allows PoW attackers to assemble attack-scale hash power in hours with no on-chain trace.

Q19. A PoS attacker holds 35% of staked ETH. The attack is detected and the correlation penalty applied. What is their capital position?
A) They recover 65% of bonded ETH B) They lose exactly 35% of their bond C) They lose nothing (penalty only applies to double-signing within a single slot) D) They lose their entire bond: $3 \times 0.35 \times \text{Bond} = 1.05 \times \text{Bond}$, capped at 100%
Answer: D – $3 \times 0.35 = 1.05$: the penalty exceeds the bond, so 100% of the bond is burned. Total capital loss.

Q20. A designer claims: "Switch to PoW because PoS has a weak subjectivity problem." What is the strongest economic counter-argument?
A) PoW has the same weak subjectivity problem as PoS, which is why both consensus mechanisms require external parties to provide trusted checkpoints for new nodes B) The weak subjectivity problem has been fully solved by Casper FFG C) Weak subjectivity needs one auditable checkpoint; PoW attack capital is rentable and recoverable, making failed attacks cheaper than a burned PoS bond D) PoS has lower energy consumption, which outweighs any trust assumption

Quiz Questions 16–20

Q16. Which dimension does this supplement add to L10's PoW/PoS consensus comparison?
A) Transaction throughput measured in confirmed transactions per second, since higher throughput implies more economic activity secured per block B) Energy consumption per transaction C) Block time and confirmation speed D) Whether attack capital can be recovered after a failed or detected attack
Answer: D – L10 covers energy/TPS/finality speed. This deck adds: capital recoverability, finality type (probabilistic vs. economic), and attack-capital assembly time.

Q17. What is the primary ongoing cost of PoS network security, as opposed to PoW?
A) ETH issuance (inflation), paid collectively by all ETH holders as dilution B) Hardware depreciation, paid by validator operators C) Gas fees, paid by transaction senders D) Electricity costs billed to each validator operator for the continuous power consumption required to run hashing hardware around the clock and compete in the proof-of-work lottery
Answer: A – PoS validators earn newly issued ETH (inflation) as rewards. PoW miners earn block rewards plus fees but require continuous energy expenditure.

Q18. Why does a would-be PoS attacker face more entry friction than a PoW attacker?
A) PoS requires the private keys of the Ethereum Foundation B) PoW attacks require cooperation from all mining pools C) PoS validator selection uses verifiable randomness that prevents any party from predicting which validators will attest to specific blocks D) PoS attack capital must be bonded on-chain (slow, visible); PoW hash power can be rented via NiceHash in hours
Answer: D – Assembling 33%+ of staked ETH is slow and publicly visible on-chain. NiceHash allows PoW attackers to assemble attack-scale hash power in hours with no on-chain trace.

Q19. A PoS attacker holds 35% of staked ETH. The attack is detected and the correlation penalty applied. What is their capital position?
A) They recover 65% of bonded ETH B) They lose exactly 35% of their bond C) They lose nothing (penalty only applies to double-signing within a single slot) D) They lose their entire bond: $3 \times 0.35 \times \text{Bond} = 1.05 \times \text{Bond}$, capped at 100%
Answer: D – $3 \times 0.35 = 1.05$: the penalty exceeds the bond, so 100% of the bond is burned. Total capital loss.

Q20. A designer claims: "Switch to PoW because PoS has a weak subjectivity problem." What is the strongest economic counter-argument?
A) PoW has the same weak subjectivity problem as PoS, which is why both consensus mechanisms require external parties to provide trusted checkpoints for new nodes B) The weak subjectivity problem has been fully solved by Casper FFG C) Weak subjectivity needs one auditable checkpoint; PoW attack capital is rentable and recoverable, making failed attacks cheaper than a burned PoS bond D) PoS has lower energy consumption, which outweighs any trust assumption
Answer: C – Weak subjectivity is bounded and auditable. PoW's recoverable attack capital (60-70% hardware resale) and rental market (NiceHash) reduce the economic cost of a failed attack relative to PoS's burned bond.