

MIBL TERM 6

Financial Risk Regulation of Financial Institutions

JPMorgan Chase and the London Whale

A Written Analysis by Risk Type: Credit, Market, Liquidity, Operational, and Compliance/Conduct Risk

Prepared using the course lecture deck, the finals study guide, and the assigned case materials:

- “JPMorgan Chase London Whale A: Risky Business” (Yale YPFS Case 2014-2a)
 - “JPMorgan and the London Whale” (INSEAD Case INS370)
- “JPMorgan Chase London Whale Z: Background and Overview” (Yale YPFS Case 2014-2z)
- U.S. Senate PSI, “JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses” (2013)
 - MIBL Term 6 lecture deck and Finals Study Guide (Prof. Camus, AIM)

Giselle Arellano

August 2026

Table of Contents

1. Introduction: Reading the Case Through Five Risk Lenses	3
2. Case Snapshot	3
3. Credit Risk	5
3.1 The Synthetic Credit Portfolio as a credit risk transfer instrument	5
3.2 Concentration risk	5
3.3 Counterparty credit risk and the collateral dispute evidence	6
3.4 Realized credit events and Expected vs. Unexpected Loss	6
4. Market Risk	7
4.1 Basis risk: the hedge that stopped matching what it hedged	7
4.2 Value-at-Risk: definition, and the January 2012 model change	7
4.3 Other market-risk metrics and their breaches	8
4.4 Duration, convexity, and non-linearity	8
4.5 Market impact: when a trading desk becomes the market	8
5. Liquidity Risk	9
5.1 Funding liquidity: never at risk, and part of the root cause	9
5.2 Asset liquidity: the real liquidity story of this case	9
5.3 Liquidity risk controls present, and the one that was missing	9
6. Operational Risk	11
6.1 People risk: unqualified and non-independent staff	11
6.2 Process risk: bypassed validation and unenforced protocols	11
6.3 Systems risk: a model with “formula and calculation errors”	11
6.4 Mismarking as a control failure	12
6.5 Capital treatment and consequences	12
7. Compliance and Conduct Risk	13
7.1 Market conduct: the CFTC manipulation finding	13
7.2 Compliance risk: regulatory reporting and the OCC relationship	13
7.3 Misleading public and investor disclosure	13
7.4 Consequences: the regulatory and legal response	14
7.5 Risk culture as the shared root of the failures	14
8. Cross-Risk Synthesis	15
8.1 The failure chain	15
8.2 Risk-type synthesis matrix	15
8.3 A brief comparative note	16
9. Conclusion	17
Sources	18

1. Introduction: Reading the Case Through Five Risk Lenses

The JPMorgan Chase “London Whale” episode of 2012 is unusual among bank-failure case studies because it cannot be explained by a single category of risk. A \$6.2 billion trading loss, a mismarked derivatives book, a manipulated internal model, misleading public disclosures, and a six-year supervisory blind spot at the Office of the Comptroller of the Currency (OCC) all emerged from one trading desk — JPMorgan’s Chief Investment Office (CIO) — within a single twelve-month window. This analysis follows the taxonomy set out in both the MIBL Term 6 lecture deck and the accompanying Finals Study Guide, examining the case through each of the five risk categories the course uses to structure bank risk management: Credit Risk, Market Risk, Liquidity Risk, Operational Risk, and Compliance/Conduct Risk. A closing synthesis then shows how the five categories interacted — because, as the lecture’s own closing slides put it, “risks rarely arrive alone.”

Each section below defines the risk category as the course defines it, applies the course’s own assessment and mitigation tools to the specific facts of the Whale trades, and identifies where JPMorgan’s controls in that category succeeded, failed, or were absent altogether. Facts are drawn primarily from the U.S. Senate Permanent Subcommittee on Investigations (PSI) report (the most detailed primary source), supplemented by the Yale YPFS case series (“Risky Business” and “Background and Overview”) and the INSEAD case (INS370), which together provide the fullest picture of the trading strategy, its risk-model mechanics, and its aftermath.

A note on classification

Several of the case’s central facts sit at the boundary between categories — credit-default-swap spread risk is simultaneously a credit exposure and a market-priced instrument; the VaR model change is simultaneously a market-risk measurement failure and an operational-risk model-governance failure; the CFTC’s manipulation finding is simultaneously a market-conduct violation and a compliance breach. Where a fact illuminates more than one risk type, it is discussed in its primary category with a cross-reference to the others, and the Section 7 synthesis matrix maps every major fact to all the risk types it touches.

2. Case Snapshot

JPMorgan Chase’s Chief Investment Office (CIO), formed in 2005 under Ina Drew to invest the bank’s roughly \$350 billion of deposits in excess of loans, also ran a Synthetic Credit Portfolio (SCP) of credit-derivative index positions authorized in 2006 as a portfolio hedge against the bank’s overall credit risk. In December 2011, CEO Jamie Dimon and CFO Douglas Braunstein instructed CIO to reduce Risk-Weighted Assets (RWA) ahead of Basel III. Rather than divest risk, lead trader Bruno Iksil built an increasingly large, complex book — growing from \$51 billion in net notional exposure at year-end 2011 to \$157 billion by the end of Q1 2012 — that lost \$100 million in January, \$69 million in February, and \$550 million in March 2012, breaching CIO’s risk limits more than 330 times in four months. Management responded by raising limits and adopting a new Value-at-Risk (VaR) model that cut reported risk by roughly half overnight, while junior trader Julien Grout, under Iksil’s direction, mismarked the book to hide losses of up to \$500 million at a time. When Bloomberg and the Wall Street Journal exposed the trades on April 6, 2012, JPMorgan defended them as a hedge; Dimon called the matter “a complete tempest in a teapot” on the April 13 earnings call. Losses reached \$2

billion by May, \$4.4 billion by June, and at least \$6.2 billion by year-end, forcing a rare earnings restatement, the termination of four traders (two later indicted for fraud), Ina Drew's resignation, and \$1.02 billion in penalties from five regulators in 2013.

3. Credit Risk

The course defines credit risk as “the risk that a borrower or counterparty will fail to meet its financial obligations as they come due, resulting in a financial loss to the lender or investor,” driven by four components: Probability of Default (PD), Exposure at Default (EAD), Loss Given Default (LGD), and Maturity (M), combining into Expected Loss = $PD \times LGD \times EAD$. The Finals Study Guide adds default, counterparty, and concentration risk as the three faces of credit risk, and names collateral, covenants, diversification, and securitization as the standard mitigants, with IFRS 9 Expected Credit Loss (ECL) and Credit Valuation Adjustment (CVA) as the modern measurement tools. Every one of these concepts appears in the Whale case — but in an inverted form, since the instrument built to manage JPMorgan’s credit risk became, itself, the case’s largest concentrated credit exposure.

3.1 The Synthetic Credit Portfolio as a credit risk transfer instrument

Credit default swaps (CDS) — the building blocks of SCP — are, by definition, a credit risk transfer/mitigation tool: the “protection buyer” pays a premium to transfer default risk on a reference entity to a “protection seller,” who assumes that credit risk in exchange for income, exactly as an insurer would (though, unlike insurance, a CDS buyer need not own the underlying exposure). CIO’s SCP was authorized in 2006 specifically to hedge JPMorgan’s own credit exposure — a legitimate, course-recognized use of derivatives as a credit risk mitigant, paralleling the lecture’s discussion of CDS as a device “to transfer risk of default from one party to another.” SCP traded baskets of single-name CDS bundled into indices — CDX.NA.IG (125 North American investment-grade names) and CDX.NA.HY (100 North American high-yield names) — meaning every SCP position carried the aggregated default risk ($PD \times LGD \times EAD$, in the course’s formula) of dozens to well over a hundred corporate borrowers at once.

By early 2012, however, SCP’s net position had flipped: rather than being predominantly long protection (a “short credit risk” hedge that pays off if borrowers default, offsetting JPMorgan’s loan-book exposure), the book became net long credit risk — meaning SCP’s risk moved in the same direction as the rest of the bank’s balance sheet rather than offsetting it. This is the single clearest credit-risk finding in the entire record: a position explicitly built to reduce JPMorgan’s aggregate credit risk instead added to it, at exactly the scale (up to \$157 billion notional) that made the addition material.

3.2 Concentration risk

The lecture’s citation of the Bank for International Settlements (BIS) defines risk concentration as “any single exposure or group of exposures with the potential to produce losses large enough to threaten a bank’s health or ability to maintain its core operations” — and SCP’s single largest position, in the CDX.NA.IG.9 series alone, exceeded \$80 billion in notional exposure to a single basket of 125 correlated corporate names by early 2012. Diversification — the Study Guide’s first-listed credit risk mitigant — is precisely what SCP lacked at this scale: unlike a diversified loan book spread across thousands of borrowers and sectors, SCP concentrated an enormous notional exposure in a small number of correlated index series (IG.9, IG.17/18, and a handful of iTraxx European series), all sharing correlated exposure to the same macro-credit cycle. When Ally Financial’s ResCap subsidiary filed for bankruptcy on February 13, 2012 — an idiosyncratic single-name credit event — SCP lost money on both its long and short legs simultaneously, direct evidence that the book’s supposed internal diversification (offsetting HY shorts against IG longs) did not actually diversify away correlated credit-event risk.

3.3 Counterparty credit risk and the collateral dispute evidence

Every CDS SCP entered into carried counterparty credit risk — the risk that the trader on the other side of the swap cannot pay out if a credit event occurs — which the lecture explicitly flags as a derivative-specific risk requiring collateral/margining discipline. This mechanism is what ultimately forced JPMorgan's hand on valuation: because CIO marked its book at prices more favorable than the true bid-ask midpoint (see Section 6.3 on operational risk), while counterparties on the other side of the same trades used honest midpoint marks, the two sides' collateral calculations diverged. Collateral disputes with ten counterparties peaked at \$690 million (the largest with Morgan Stanley) — a textbook instance of collateral, the primary credit-risk mitigant for derivatives exposure, functioning exactly as designed: an outside counterparty's independent valuation caught an internal mismarking problem that JPMorgan's own Valuation Control Group had not stopped. This is also a direct, real-world illustration of Basel's Pillar 3 "market discipline" concept from the lecture — external parties with their own money at stake enforced a check that internal controls failed to provide.

The Basel III capital framework's CVA (Credit Valuation Adjustment) capital charge — cited in the Finals Study Guide as a modern credit-risk tool — is directly relevant to why this position existed in the form it did: derivatives exposure under Basel III draws capital charges for both counterparty credit risk and CVA risk on top of ordinary market-risk charges, which is precisely the RWA increase Dimon and Braunstein were trying to shrink in December 2011. In other words, the credit-risk capital regime itself is the proximate cause of the strategy that produced the loss (see Section 4 and Section 5.4 for the market-risk and Basel-driven dimensions of this same fact).

3.4 Realized credit events and Expected vs. Unexpected Loss

Two actual corporate bankruptcies bookend the strategy's formation: American Airlines' November 29, 2011 bankruptcy netted SCP an estimated \$400–550 million gain on protection held since summer 2011 — a realized credit event working in JPMorgan's favor, consistent with Expected Loss modeling on a diversified short-risk book. But Eastman Kodak's January 19, 2012 bankruptcy cost SCP a \$50 million loss, because CIO had let offsetting purchased protection lapse in December 2011 specifically to help reduce RWA (purchased protection counts as a balance-sheet asset that inflates RWA) — meaning a deliberate, RWA-driven reduction in credit protection directly caused an avoidable credit loss weeks later. This is the course's Unexpected Loss concept made concrete: JPMorgan's own capital-optimization decisions widened the tail of outcomes the desk was exposed to, and Kodak fell squarely into that widened tail.

Credit risk verdict

SCP was authorized, measured, and defended publicly as a credit-risk mitigation tool — but by Q1 2012 it had become an under-diversified, poorly documented credit exposure in its own right, one whose size (>\$80B in a single correlated index series) met the BIS definition of a concentration risk capable of threatening the bank's health, and whose valuation discipline was ultimately enforced not by internal controls but by counterparties disputing collateral calls.

4. Market Risk

The course defines market risk as “the possibility of financial losses due to adverse movements in market prices or rates,” arising from interest rate, equity, FX, and commodity exposure, and — most relevant here — measured through standard deviation, beta, duration, and Value-at-Risk (VaR). Market risk is explicitly framed in the lecture as systematic: it “affects all participants in the market and cannot be eliminated through diversification.” Credit-spread risk (the price risk of CDS premiums moving as credit sentiment shifts) sits at the intersection of credit and market risk — SCP’s positions were, in market-risk terms, bets on the relative movement of high-yield versus investment-grade credit spreads, which is why VaR, not a pure credit model, was CIO’s primary risk-measurement tool for the book. This section is the deepest and most heavily documented part of the case.

4.1 Basis risk: the hedge that stopped matching what it hedged

The lecture’s treatment of “Risk of Using Derivatives Improperly” names basis risk explicitly: “when the derivative doesn’t perfectly match the hedged item.” This is the cleanest one-line description of what happened to SCP between 2011 and 2012. Starting in summer 2011, Iksil began buying CDX.NA.HY protection while selling CDX.NA.IG protection to fund the premium cost — a strategy with an increasingly loose connection to any specific, documented balance-sheet exposure. Senate PSI Finding #2 states this directly: JPMorgan never identified what assets SCP was hedging, never tested the hedge’s effectiveness, and never demonstrated that it lowered (rather than raised) the bank’s risk. A position that cannot show it moves opposite to what it claims to hedge has, by definition, unmanaged basis risk — and by March 2012 SCP had flipped to net long credit risk, moving in the same direction as the rest of the bank rather than against it.

4.2 Value-at-Risk: definition, and the January 2012 model change

VaR “estimates maximum expected loss over a defined time period at a given confidence level” — the course’s worked example computes a 95%-confidence 1-day VaR from historical return percentiles. The lecture’s own “Criticisms of VaR During the Crisis” slide states that “internal models at some investment banks were designed to minimize regulatory capital, not truly measure risk,” and that “traders and banks could manipulate inputs... to artificially reduce VaR” — a description that fits CIO’s VaR model change almost exactly.

- **What happened:** A new VaR model, built by a quantitative analyst (Patrick Hagan) who had never designed a VaR model before, was adopted in late January 2012 — while CIO was actively in breach of its own and the bank-wide VaR limit — without completing the bank’s standard practice of running old and new models concurrently for validation, and without OCC approval. The new model cut CIO’s reported 95% VaR by roughly 44–50% overnight.
- **Effect on disclosure:** Q1 2012 average CIO VaR was first disclosed (April 13, 2012) at \$67 million; once the bank reverted to the old model, it was restated (May 10, 2012) to \$129 million — essentially double, and roughly double the 2011 average of \$60–61 million.
- **Proof of failure:** On April 10, 2012, SCP lost \$415 million in a single day — more than six times the ~\$67 million VaR figure supposedly bounding a 1-in-20 worst day. The model’s own developer later told the Senate the model was simply “wrong.”

This connects to the Study Guide’s Learning Objective 3 tools directly: VaR, Expected Shortfall, and Stress Testing are listed as the core market-risk toolkit, and the course lecture notes that under the post-crisis Fundamental Review of the Trading Book (FRTB), VaR is being replaced by Expected Shortfall at 97.5% confidence precisely because VaR — as this case demonstrates — understates tail risk and can be gamed at the model level. Had SCP been measured on an Expected Shortfall basis (which by construction averages the losses beyond the confidence threshold, rather than reporting only the threshold itself), the magnitude of the tail the book was exposed to would have been far harder to mask through a single model recalibration.

4.3 Other market-risk metrics and their breaches

Beyond VaR, CIO tracked CS01/CSBPV (Credit Spread Basis-Point Value — dollar sensitivity to a 1bp spread move, a linear “credit duration” measure that ignores convexity) and CSW10% (a stress test for a simultaneous 10% widening of every position’s spread). Between January 1 and April 30, 2012, CIO’s risk limits and advisories were breached more than 330 times: CS01 was exceeded by 100% in January, 270% in early February, and over 1,000% by mid-April; CSW10% was breached repeatedly in March, with the March 22 breach (on a single-day sale of \$10 billion+ notional) triggering Ina Drew’s March 23 “put phones down” trading halt. A February 22, 2012 run of CIO’s Comprehensive Risk Measure (CRM) — a stress-scenario tool — projected a potential \$6.3 billion annual loss; a CIO risk manager dismissed the figure as “garbage” rather than escalating it. All of these are stop-loss and stress-testing tools the Study Guide lists as standard market-risk mitigants — present at CIO in name, but institutionally ignored rather than acted upon.

4.4 Duration, convexity, and non-linearity

The lecture’s duration/convexity module notes that bond price sensitivity is convex, not linear, and flags — in the instructor’s own annotation — that “Value at risk[,] convex, doesn’t go down this way... NON LINEAR.” CS01, a linear sensitivity measure, is structurally unable to capture this convexity, which is exactly why CIO also tracked CSW10% as a second-order stress measure. But a position built from CDS index tranches carries genuine convexity and correlation-driven non-linearity (echoing the lecture’s Basel III diagram showing that “the impact of correlation results in a fatter-tailed distribution” than a simple credit model assumes) — meaning CIO’s reliance on VaR and CS01 as primary metrics, with CSW10% treated as a secondary check that was itself frequently breached and ignored, left the desk structurally under-equipped to see the size of tail losses building in the book.

4.5 Market impact: when a trading desk becomes the market

Market risk models generally assume a position can be priced (and, if needed, exited) at prevailing market prices. That assumption failed here on a scale rarely seen outside a genuine market squeeze: on February 29, 2012, Iksil sold protection on over \$7 billion notional of CDX.NA.IG.9 in a single day — more than 90% of that day’s total market volume in the instrument. INSEAD Case INS370’s data shows CIO’s share of total street trading volume reaching 48–49% across several index/tenor combinations between January and April 2012. Once rival traders recognized JPMorgan was behind the outsized, one-directional selling, they began trading against the known position, and — per Case A’s account — could buy protection on the underlying single-name companies more cheaply than buying the JPMorgan-distorted index itself, a clean signal that the market had priced in JPMorgan’s footprint. This dimension of market risk — the price impact of one’s own trading — is largely invisible to a standard VaR model, which typically assumes the trader is a price-taker, not a price-setter.

Market risk verdict

SCP's market risk failure was compounded, not singular: a basis-risk drift away from any documented hedge target, a VaR model deliberately recalibrated downward while in active breach, secondary stress metrics that were breached hundreds of times and overridden rather than enforced, and a position so large it began to set, rather than take, market prices in the instruments it traded.

5. Liquidity Risk

The course's central liquidity distinction is between Funding Liquidity Risk ("an institution's inability to meet its liabilities and obligations as they come due without incurring unacceptable losses") and Asset Liquidity Risk ("an asset position cannot easily be unwound or offset at short notice without significantly influencing the market price, because of inadequate market depth or market disruption"), summarized in the course's recurring line: "you can be solvent and still die of illiquidity." The Whale case is unusual among the course's liquidity case studies (Northern Rock 2007 and Urban Bank 2000, per the Finals Study Guide) in that JPMorgan itself never came close to a funding liquidity crisis — the liquidity story here plays out entirely on the asset side, and at the level of a single trading book rather than the whole bank.

5.1 Funding liquidity: never at risk, and part of the root cause

JPMorgan held \$1.128 trillion in deposits against only \$724 billion in loans at year-end 2011 — a roughly \$400 billion structural excess of stable, low-cost funding. CIO existed specifically to invest this excess, and the bulk of it sat in a conservative, AA+-rated \$350 billion bond portfolio, not in SCP. This abundance of funding liquidity is, in a real sense, part of the root cause of the case: it is precisely because CIO had such a large, low-scrutiny pool of capital to deploy — nearly double JPMorgan's \$184 billion in total stockholders' equity — that a comparatively small trading unit could build a \$157 billion notional derivatives book without triggering the funding-side alarms that would accompany a bank under liquidity stress. JPMorgan's Basel III liquidity metrics (the Liquidity Coverage Ratio and Net Stable Funding Ratio, per the course's liquidity module) were never in question during the episode — which is itself an important lesson: liquidity ratios calculated at the whole-bank level provide no visibility into concentrated illiquidity risk building inside a single desk.

5.2 Asset liquidity: the real liquidity story of this case

SCP's core liquidity problem was textbook Asset Liquidity Risk. By March 31, 2012, the portfolio's net notional had roughly tripled to \$157 billion, concentrated in index series large enough that CIO's own trading, on certain days, represented 90%+ of total market volume. An internal JPMorgan analysis in late April 2012 found the book's largest position would take 10–15 trading days to unwind even trading at 100% of normal daily volume — meaning the position could not be exited without moving the market against itself for two full weeks even under optimistic assumptions. This is precisely the lecture's definition of asset illiquidity in action: "inadequate market depth" relative to position size. The lecture's discussion of asset fungibility notes that even highly fungible instruments can become illiquid "in stressed or disrupted markets" due to reduced market depth, volatility, and widened bid-ask spreads — and CIO's own trading was itself a major contributor to that disruption once the position size became known to counterparties.

Once the April 6, 2012 media reports named the trades, this dynamic accelerated: rival hedge funds and banks, having identified JPMorgan's position, traded against it — buying protection on individual underlying names more cheaply than through the JPMorgan-distorted index, and generally positioning to profit as JPMorgan tried to unwind. This is a liquidity spiral in miniature: increased attention and reduced counterparty willingness to trade at pre-disclosure prices widened effective bid-ask spreads on exactly the instrument JPMorgan most needed to exit, raising the realized cost of unwinding well above what the original \$500–600 million unwind estimates (from December 2011/January 2012) had assumed.

5.3 Liquidity risk controls that were present, and the one that was missing

At the firm level, JPMorgan's liquidity risk management — contingency funding plans, deposit stability monitoring, LCR/NSFR compliance (the Study Guide's named liquidity mitigants) — functioned exactly as intended and were never tested by this episode, because the loss never threatened the bank's ability to meet its obligations. What was missing was a control specific to concentrated position-level asset liquidity: CIO lacked concentration limits on its derivatives positions, a standard control that existed at JPMorgan's own Investment Bank for economically identical instruments. A concentration limit — capping notional exposure to a single index series or tenor as a percentage of typical market volume — is the natural liquidity-risk analog to the credit-risk concentration limits discussed in Section 3.2, and its absence at CIO meant no internal control forced a halt before the position had already grown large enough to distort the market it traded in.

Liquidity risk verdict

JPMorgan Chase, the institution, was never illiquid — this case offers no analog to Northern Rock's 2007 deposit run or a Lehman-style funding freeze. The liquidity failure here is narrower and more specific: a single trading book grew, through the absence of a concentration limit, to a size where it could not be liquidated without moving the market against itself — proof that whole-bank liquidity ratios (LCR/NSFR) do not, by themselves, capture desk-level asset liquidity risk.

6. Operational Risk

Basel II's definition, used verbatim in the course lecture, is “the risk of loss resulting from inadequate or failed internal processes, people, and systems or from external events.” The lecture is explicit that operational risk, unlike credit, market, and liquidity risk, is not primarily financial in nature but concerns “the way an institution operates,” and can produce indirect losses — reputational damage, regulatory penalties — beyond the immediate dollar loss. Notably, the course's own lecture deck uses JPMorgan's \$6.2 billion loss as its explicit illustrative example of an operational-risk event (shown alongside Wells Fargo), which signals that the instructor's own frame for this case leads with operational risk, not market risk, as the root category — this analysis treats operational risk as the connective tissue linking the market-risk symptoms (Section 4) to the compliance consequences (Section 7).

6.1 People risk: unqualified and non-independent staff in critical control roles

The Finals Study Guide names “segregation of duties” as the primary operational-risk mitigant, and its violation is one of the clearest facts in the entire case. Julien Grout, a junior trader who reported directly to Iksil, was responsible for marking the book's daily value under Iksil's direction — meaning the person entering the SCP's official valuation had no independence from the person whose P&L that valuation determined. CIO had no dedicated Chief Risk Officer at all until January 2012, when Irvin Goldman — a capital-markets veteran with no prior risk-management experience — was appointed, having been nominated by his own brother-in-law, the outgoing bank-wide CRO Barry Zubrow. The new VaR model itself was built by quantitative analyst Patrick Hagan, who had never designed a VaR model before. Each of these is a “people” failure in the Basel II sense: critical control functions staffed by individuals lacking either the independence or the expertise the role required.

6.2 Process risk: bypassed model validation and unenforced risk-limit protocols

JPMorgan's standard process for adopting a new risk model required running the new and old models concurrently for several months before approval, and required OCC sign-off for a model applied to a portfolio of SCP's significance. Both steps were skipped in the rush to resolve the January 2012 VaR breach. The Study Guide lists back-testing and Key Risk Indicators (KRIs) as standard operational-risk monitoring tools — CIO's risk metrics (VaR, CS01, CSW10%, the CRM) functioned as exactly this kind of KRI system, and the process for responding to a KRI breach (330+ times between January and April 2012) should have been escalation and remediation. Instead, per the Senate Report, breaches were “routinely reported” up the chain but “largely ignored or ended by raising the relevant limit” — a process failure in how breaches were handled, not merely a failure to detect them.

6.3 Systems risk: a model with “formula and calculation errors”

The new VaR model was later found to be “improperly implemented,” requiring error-prone manual data entry and containing outright formula and calculation errors — a direct instance of the Basel II definition's “failed... systems” clause. This is a useful case for distinguishing operational risk from market risk analytically: the fact that the new model produced a lower VaR number is a market-risk-measurement problem (Section 4.2), but the fact that the model contained calculation errors and required manual, error-prone data entry is a systems/process problem regardless of what number the model happened to produce

— even a model calibrated to be conservative rather than aggressive would still have been an operational-risk failure if built this way.

6.4 Mismarking as a control failure, not just a market-risk symptom

The mismarking practice — CIO marking SCP positions at prices more favorable than the bid-ask midpoint from late January through March 2012, hiding as much as \$432–512 million in losses at various points — is properly understood as an operational-risk (internal control) failure with market-risk and compliance consequences, not a market-risk failure in itself. The valuation control process that should have caught it — an independent Valuation Control Group (VCG) and Controller’s office review — reviewed the marks and raised no objection at the time, illustrating a control that existed on paper but did not function independently in practice. The practice ended only when Deputy Chief Risk Officer Ashley Bacon ordered CIO to use the same independent pricing service as the Investment Bank — an external, non-CIO intervention, not an internal control catching its own failure.

6.5 Capital treatment and consequences

Basel’s operational risk capital charge — under the framework the Study Guide summarizes as evolving from a flat percentage of gross income (Basel II) to the Basel III/IV Standardized Approach’s Business Indicator combined with an Internal Loss Multiplier — means realized operational losses of this scale directly raise a bank’s future required capital: “large banks with significant losses will hold more capital due to the internal loss multiplier,” in the course’s own words. JPMorgan’s \$6.2 billion loss and the associated \$1.02 billion in regulatory penalties therefore carry a forward-looking capital consequence beyond the immediate dollar cost — a mechanism the lecture frames as Basel’s way of making operational risk “count” in capital terms even though, unlike credit or market risk, it cannot be modeled with the same statistical precision.

Operational risk verdict

Read through the Basel II definition, nearly every proximate cause of the loss — an under-independent junior trader marking his own supervisor’s book, a first-time model-builder’s VaR model with calculation errors, a validation process bypassed under time pressure, and risk-limit breaches processed as paperwork rather than triggers for action — is an operational risk failure. This is consistent with the course’s own choice to use this case as its illustrative operational-risk example.

7. Compliance and Conduct Risk

The course draws a specific distinction between two related but separate questions: Compliance Risk asks “are we following the rules?” — the risk of legal or regulatory sanction from failing to comply with laws, regulations, or internal policy. Conduct Risk asks “are we treating customers, markets, and counterparties fairly?” — noting explicitly that “an action may be technically legal but still create conduct risk if it is unfair, misleading, abusive, or inconsistent with customer interests,” with market misconduct named as a specific example. The Whale case produces clear violations of both, plus a third dimension — misleading public disclosure — that sits closest to securities-law compliance.

7.1 Market conduct: the CFTC manipulation finding

On February 29, 2012, Iksil sold protection on over \$7 billion notional of CDX.NA.IG.9 — more than 90% of that day’s total market volume — large enough to single-handedly suppress the market price of IG.9 protection. The Commodity Futures Trading Commission (CFTC) later determined this trade violated the Commodity Exchange Act’s manipulation provisions, and JPMorgan admitted, as part of its October 2013 settlement, that CIO traders “acted recklessly.” This is conduct risk in the course’s precise sense: the trade was, in form, an ordinary derivatives sale, but its scale relative to the market made it unfair to other market participants and abusive of JPMorgan’s outsized market footprint — “technically” a normal trade, but a conduct violation in substance.

7.2 Compliance risk: regulatory reporting and the OCC relationship

JPMorgan never disclosed SCP’s existence to the OCC, its primary regulator, in any regular report before the trades became public — despite SCP growing tenfold in 2011 alone. In February–March 2012, the bank stopped delivering a required CIO management report to the OCC without explanation; the OCC did not notice the gap until after the April 6 media exposure. When directly asked whether CIO had mismarked the book, the bank’s Chief Risk Officer “initially denied it.” Each of these is a compliance-risk violation in the narrowest sense: a failure to meet specific, identifiable regulatory reporting and candor obligations, independent of whether the underlying trading itself was against the rules.

7.3 Misleading public and investor disclosure

On the April 13, 2012 earnings call, CFO Douglas Braunstein stated that SCP was “fully transparent to the regulators,” that positions were “put on pursuant to the risk management at the firm-wide level,” that decisions were “made on a very long-term basis,” and that the strategy would be “allowable under the Volcker Rule.” The Senate PSI report found each of these claims false or seriously misleading: SCP had never been disclosed to the OCC in any regular report; firm-wide risk managers, including CRO John Hogan, had no role in the positions; an internal CIO presentation described SCP’s horizon as “short”; and the bank had, in a February 2012 letter, itself expressed concern that similar activity could be prohibited under the Volcker Rule. CEO Dimon’s characterization of the matter as “a complete tempest in a teapot” was made, per the Senate Report, while he already knew of three straight months of escalating losses, the book’s tripling in size, its flip to net-long credit risk, and the March 23 trading halt. The Senate Report’s conclusion — that these statements were “fictions irreconcilable with the bank’s obligation to provide material information to its investors” — frames this squarely as a potential Rule 10b-5/Securities Act §17(a) issue, i.e., compliance risk in the disclosure-law sense, not merely a public-relations failure.

7.4 Consequences: the regulatory and legal response

Date	Regulator / Authority	Outcome
Sept. 19, 2013	US Federal Reserve, OCC, SEC; UK FCA	\$920 million combined; JPMorgan required to admit wrongdoing in certain instances.
Oct. 16, 2013	US CFTC	\$100 million civil penalty; JPMorgan admitted CIO traders “acted recklessly” and that the Feb. 29, 2012 sale violated manipulation provisions of the Commodity Exchange Act.
Aug. 14, 2013	SEC / DOJ	Martin-Artajo and Grout charged with fraud (in absentia); Iksil entered a non-prosecution agreement.
Jan. 14, 2013	OCC	Formal Cease-and-Desist Order, following six 2012 Supervisory Letters citing 20 “Matters Requiring Attention.”

Total combined regulatory penalties reached approximately \$1.02 billion — an unusually large and unusually candid (wrongdoing-admitting) package, reflecting the multiplicity of compliance and conduct failures found (market manipulation, disclosure violations, and regulatory-candor breakdowns) rather than a single infraction.

7.5 Risk culture as the shared root of the compliance/conduct failures

The course’s Senior Management Failure Patterns list — “frames issues as ‘isolated incidents’ when they are actually systemic,” “treats compliance as a cost center rather than a core control,” “ignores weak signals from audit, compliance, or customer complaints” — maps directly onto the compliance and conduct facts above: Dimon’s “tempest in a teapot” remark is the clearest instance of framing a systemic problem as isolated; the decision to keep growing SCP rather than pay the \$500–600 million unwind cost is the clearest instance of treating a control cost as expendable; and the dismissal of the \$6.3 billion CRM stress result as “garbage” is the clearest instance of ignoring a weak (but real) signal. This is the same risk-culture root cause identified in the previously-prepared thematic study guide’s governance analysis — Section 7 here simply traces its compliance and conduct consequences specifically.

Compliance/conduct risk verdict

The case produces violations on all three compliance/conduct fronts the course distinguishes: market conduct (the CFTC manipulation finding), regulatory compliance (undisclosed reporting to the OCC, denied mismarking when asked), and disclosure compliance (the April 13 earnings-call misstatements). The \$1.02 billion in penalties and the rare requirement to admit wrongdoing reflect regulators treating this as a multi-front compliance failure, not an isolated trading loss.

8. Cross-Risk Synthesis: Why the Categories Cannot Be Read in Isolation

The course's closing lecture slides make a point worth using as this analysis's organizing thesis: "The 2007–2008 Financial Crisis: Proof That Risks Don't Fail Alone," tracing a chain from credit risk (subprime lending) through market risk (securitization spreading exposure) to liquidity risk (funding markets freezing) to a governance failure (exposures mismeasured). The Whale case follows an almost identical chain, compressed into a single trading desk over four months instead of a financial system over several years.

8.1 The failure chain

1. A Basel III capital rule change (credit-risk capital treatment of derivatives) created pressure to cut Risk-Weighted Assets cheaply rather than by simply reducing exposure.
2. The resulting trading strategy generated a concentrated credit-risk position (Section 3) that was also a compounding market-risk position (Section 4), too large to exit without moving the market against itself (Section 5 — asset liquidity risk).
3. As losses emerged, operational-risk failures — an unvalidated VaR model, non-independent marking, bypassed escalation processes (Section 6) — hid the true scale of both the market and credit risk from risk reports, management, and the board.
4. Once the position became public, the same underlying facts produced compliance and conduct violations (Section 7): a market-manipulation finding, undisclosed regulatory reporting, and misleading public statements — because the operational failure to measure risk accurately fed directly into public statements that were, in substance, false.
5. Capital (Basel Pillar I) and liquidity (Basel Pillar II) at the whole-bank level were never threatened — the entire chain played out within a single business unit without endangering JPMorgan's solvency or its ability to meet obligations, which is precisely why the failure is best characterized, as the previously-prepared thematic study guide concluded, as a Pillar III governance failure that happened to manifest across all five operational risk categories simultaneously.

8.2 Risk-type synthesis matrix

Risk Type	Primary Manifestation in the Case	Course Tool That Should Have Caught It	Did the Tool Work?
Credit	SCP flipped from a diversifying credit hedge to a \$80B+ concentrated, correlated credit exposure (IG.9 series)	Concentration limits; diversification; CVA capital charge	No internal concentration limit existed at CIO; caught only via counterparty collateral disputes (\$690M) — external, not internal, discipline.
Market	VaR model manipulated to halve reported risk while in active breach; basis risk from an undocumented hedge target	VaR / Expected Shortfall; CS01; CSW10%; stress testing	Metrics existed and were breached 330+ times, but breaches were "largely ignored or ended by raising the

Risk Type	Primary Manifestation in the Case	Course Tool That Should Have Caught It	Did the Tool Work?
			relevant limit” rather than acted upon.
Liquidity	\$157B position could not be unwound without moving the market; 10–15 days to exit largest position at full market volume	Concentration limits; asset liquidity monitoring	No position-level concentration limit existed at CIO (unlike the Investment Bank for the same instruments).
Operational	Non-independent marking (Grout under Iksil); unvalidated, error-containing VaR model; bypassed model-governance process	Segregation of duties; model validation; KRIs/back-testing	Segregation of duties absent; validation process explicitly bypassed “to end the breach.”
Compliance/Conduct	CFTC-confirmed market manipulation; non-disclosure to OCC; misleading April 13 earnings-call statements	Regulatory reporting protocols; disclosure controls (Reg FD / Rule 10b-5)	Failed on all fronts — \$1.02B in penalties and a rare admission of wrongdoing followed.

8.3 A brief comparative note

The Finals Study Guide pairs each risk category with a different classic case — Barings Bank (1995) for market/operational risk, Urban Bank (2000) for credit/liquidity risk, Northern Rock (2007) for liquidity risk — and the comparison is instructive. Barings is the closest analog: a single trading desk (Nick Leeson in Singapore), operating with inadequate segregation of duties, hid losses through unauthorized accounts until the losses became large enough to destroy the bank outright. The Whale case shares Barings’ core operational-risk mechanism (a trader effectively marking or controlling his own book) but differs in outcome only because JPMorgan’s \$2.3 trillion balance sheet could absorb a \$6.2 billion loss that would have been fatal to a bank the size of Barings — a useful illustration that the same governance failure can be either catastrophic or merely very costly, depending entirely on the capital base absorbing it, not on the quality of the risk management that allowed it to happen.

9. Conclusion

Viewed risk-type by risk-type, the JPMorgan Chase London Whale case resists any single-cause explanation. It is a credit risk case, because a book meant to hedge default risk became a concentrated, correlated credit exposure in its own right. It is a market risk case, because the desk's own VaR model was recalibrated to hide a real and growing risk. It is a liquidity risk case, because the position outgrew the market it traded in. It is an operational risk case — and, per the course's own framing, arguably the central one — because nearly every other failure traces back to inadequate segregation of duties, bypassed validation processes, and a risk-model system built with actual calculation errors. And it is a compliance and conduct risk case, because the same facts that were operationally mismeasured were then, in substance, misrepresented to regulators, counterparties, and investors.

The five-category framework the course uses to teach bank risk management is not simply an academic organizing device — applied rigorously to this case, it demonstrates precisely why banks are required to manage all five categories in parallel, with independent, empowered control functions in each: a bank that is well-capitalized (credit) and well-measured on a standalone basis (market) can still suffer catastrophic reputational and regulatory harm if its liquidity is concentrated at the desk level, its operational controls are compromised, and its compliance/disclosure functions do not catch what the other four categories missed. JPMorgan Chase remained solvent and liquid throughout 2012 — proof that capital and liquidity regulation, on their own, are necessary but not sufficient, and that governance and control quality across credit, market, operational, and compliance functions is what ultimately determines whether a risk-taking decision remains a manageable trading loss or becomes a systemic threat.

Sources

Asian Institute of Management, MIBL Term 6, “Financial Risk Regulation of Financial Institutions” (course lecture slides, Sessions 1–18), and accompanying Finals Study Guide, instructor J.V. Camus.

Zeissler, Arwin G.; Ikeda, Daisuke; Metrick, Andrew. “JPMorgan Chase London Whale A: Risky Business.” *Journal of Financial Crises*, Vol. 1, Iss. 2 (2019), pp. 40–59. Yale Program on Financial Stability Case Study 2014-2a-V1.

Chen, Andrew (under supervision of Claudia Zeisberger). “JPMorgan and the London Whale.” INSEAD Case INS370, ref. 03/2014-6003 (2014).

Zeissler, Arwin G.; Bennett, Rosalind; Metrick, Andrew. “JPMorgan Chase London Whale Z: Background and Overview.” *Journal of Financial Crises*, Vol. 1, Iss. 4 (2019), pp. 23–44. Yale Program on Financial Stability Case Study 2014-2z-V1.

United States Senate, Permanent Subcommittee on Investigations, Committee on Homeland Security and Governmental Affairs. “JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses.” Report, March 15, 2013.