Executive Summary

Introduction

Pampered Pets is preparing to expand its operations by introducing an international supply chain and automated warehouses. The planned digitalisation has attracted new high-profile customers, including HRH the King and Prince Albert II of Monaco. While this provides clear reputational and marketing benefits, it also raises expectations and reduces tolerance for product quality or availability failures.

The planned operational and supply chain changes are intended to support growth and efficiency. At the same time, they introduce new dependencies that may affect product quality and supply reliability. When key activities become dependent on external systems and technologies, operational risk increases (Sørensen, 2018).

This executive summary examines how digitalisation and internationalisation affect quality control and supply chain stability, with particular attention to the expectations associated with high-profile customers. A Monte Carlo simulation is used to estimate the probability of quality degradation and stockout events under uncertainty. In this context, availability is modelled as the likelihood of a stockout, reflecting the availability element of the CIA triad. Monte Carlo methods are well-suited to this type of analysis because they explicitly represent uncertainty rather than relying on fixed assumptions (Metropolis, 1987; Aven and Thekdi, 2025).

Risk to quality and Supply Chain

Digitalisation changes how production and quality control are carried out. Introducing automated processes and spreading production across multiple locations creates several related risks. These include the risk of reduced product quality due to differences in process standards, the risk of supply chain disruption caused by variable lead times and logistics delays, and the risk of temporary availability issues, such as stockouts.

Automation reduces manual handling but introduces new points of failure if calibration, configuration, or data inputs are incorrect. When production is distributed across locations, even minor process differences can affect consistency and quality (Fahimnia et al., 2019). In the model, this is reflected through a defect probability that varies across scenarios.

An international supply chain adds further uncertainty to supply availability. Customs procedures, tariffs, logistics constraints, and exchange-rate volatility all affect the predictability of supply flows. Lead times may change due to transport delays or regional disruption, while demand can fluctuate from week to week. Together, these factors make it harder to maintain the correct stock level at the right time. In the model, this uncertainty is captured through variable demand and lead times, with stockouts used as the measure of availability risk.

Taken together, these risks define the model parameters. Quality risk is represented through a variable defect probability, while supply chain uncertainty is captured through variable demand and lead times. A Monte Carlo simulation is then used to estimate the probability of quality degradation and stockout events under these conditions.

Quantitative Methodology (Monte Carlo Simulation)

Lead Time Uncertainty

Lead times were modelled as uncertain rather than fixed, reflecting the structural gap between the time required to source and deliver products and customers' expectations of availability (Christopher, 2016). Empirical studies of global and digitally enabled supply chains show that delivery times may vary by several weeks and can increase by a factor of two to four due to logistical and operational uncertainty (Ivanov and Dolgui, 2021). Lead time uncertainty was therefore represented using a triangular distribution with minimum, most likely, and maximum values.

Quality Defect Rates

Quality risk was modelled using a probabilistic defect rate rather than a fixed value. Empirical manufacturing studies report observed defect rates of 2–5% under normal operating conditions, even in controlled production environments, indicating the typical magnitude of quality variation in automated production processes (Fadli, 2024, p. 79). Based on this evidence, a conservative baseline defect probability of 1% was assumed, with quality degradation defined as scenarios where defect rates exceed 1.5%.

Demand Uncertainty

Demand was modelled as a stochastic variable. Demand variability was measured using the coefficient of variation (COV), defined as the ratio of the standard deviation to the mean demand. Weekly demand was assumed to follow a normal distribution with a mean of 180 units and a standard deviation of 30 units, producing a coefficient of variation of approximately 17%. This falls within the range associated with moderate demand variability in inventory research, where values between 10–20% are classified as moderate (Demiray Kırmızı, Ceylan and Bulkan, 2024, Table 6, p. 10).

The simulation was based on fixed input values reflecting these assumptions. Weekly demand was modelled over 52 weeks. Lead times were defined with a minimum of 1 week, a most likely value of 2 weeks, and a maximum of 5 weeks. Initial inventory was set at 800 units, with reordering triggered at 400 units and replenishment orders of 600 units. The simulation was run for 10,000 iterations to estimate the probability of quality degradation and stockout events under uncertainty.

Results and Business Impact

The Monte Carlo simulation was used to estimate the probability of quality degradation and stockout events over a 52-week period, accounting for uncertainty in demand, lead times, and product quality. The results indicate how often these risks may occur over a one-year period, consistent with probability-based risk measures used in recent risk modelling literature (Aijaz and Nazir, 2024). These estimates are indicative rather than exact and reflect outcomes under the model's simplifying assumptions.

Results from the simulation

Table 1 shows the estimated probabilities of the key risk events. Quality degradation occurs when the defect probability for the year exceeds the chosen threshold. A joint event is a scenario in which both quality degradation and at least one stockout occur within the same annual period. It should be interpreted as co-occurrence rather than causality.

Table 1: Estimated event probabilities (author's own Monte Carlo simulation)

Events	Probability
Quality degradation	16.13%
At least one stockout	99.99%
Joint quality and stockout	16.13%
Neither event	0.01 %

Figure 1 visualises the probabilities of quality degradation and joint events. The two probabilities are identical, indicating that all scenarios in which quality degradation occurs and at least one stockout occurs during the same annual period. This suggests that quality issues tend to occur alongside broader supply chain disruption rather than in isolation.

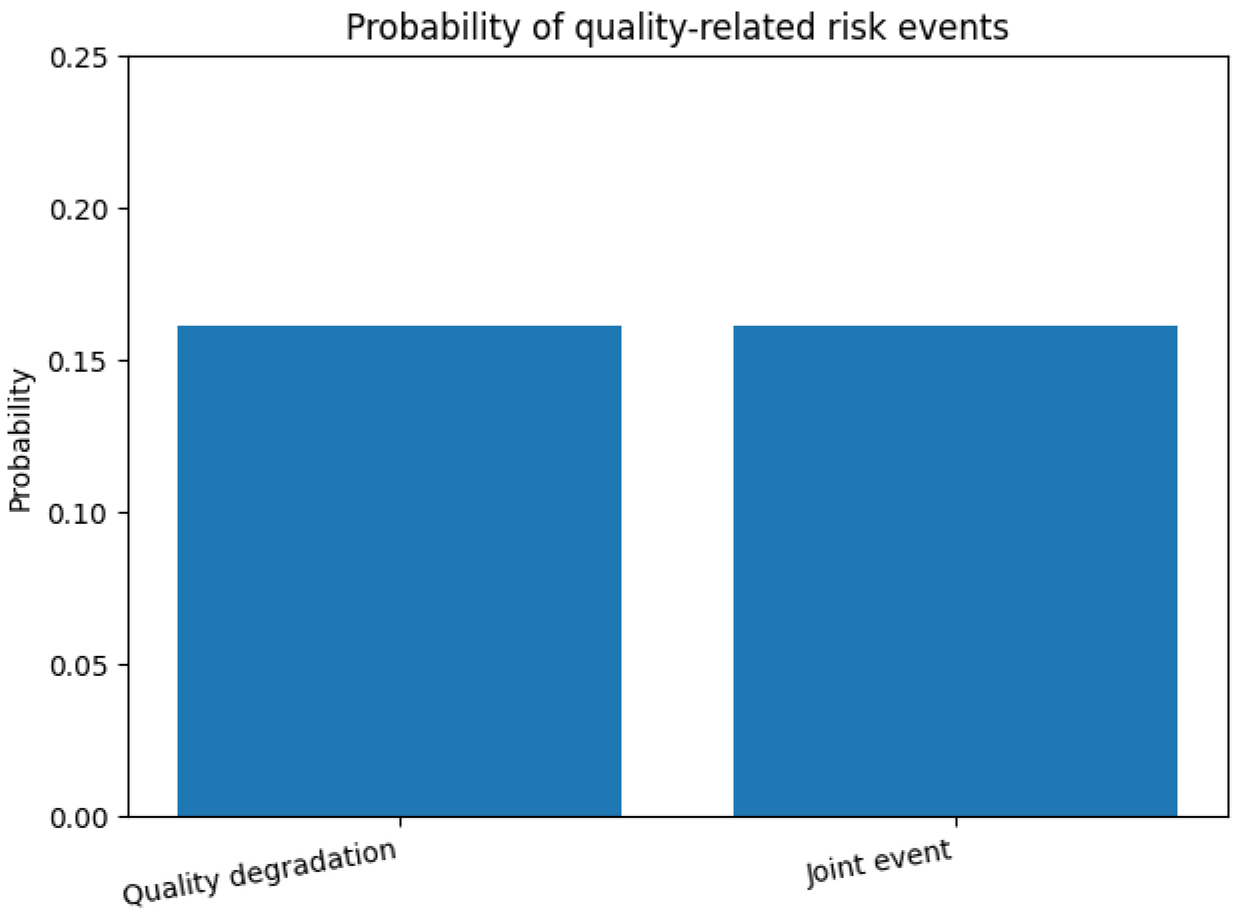

Figure 1: Annual probabilities of quality degradation and joint events (author’s own simulation)

Figure 2 shows the distribution of stockout weeks per year across simulations. While stockouts occur in most simulated years, the distribution indicates that their impact is typically limited in duration. Most scenarios experience only a small number of stockout weeks, while more prolonged shortages are relatively rare. This suggests that availability issues are frequent but usually short-lived, reflecting a persistent operational risk rather than a structural one.

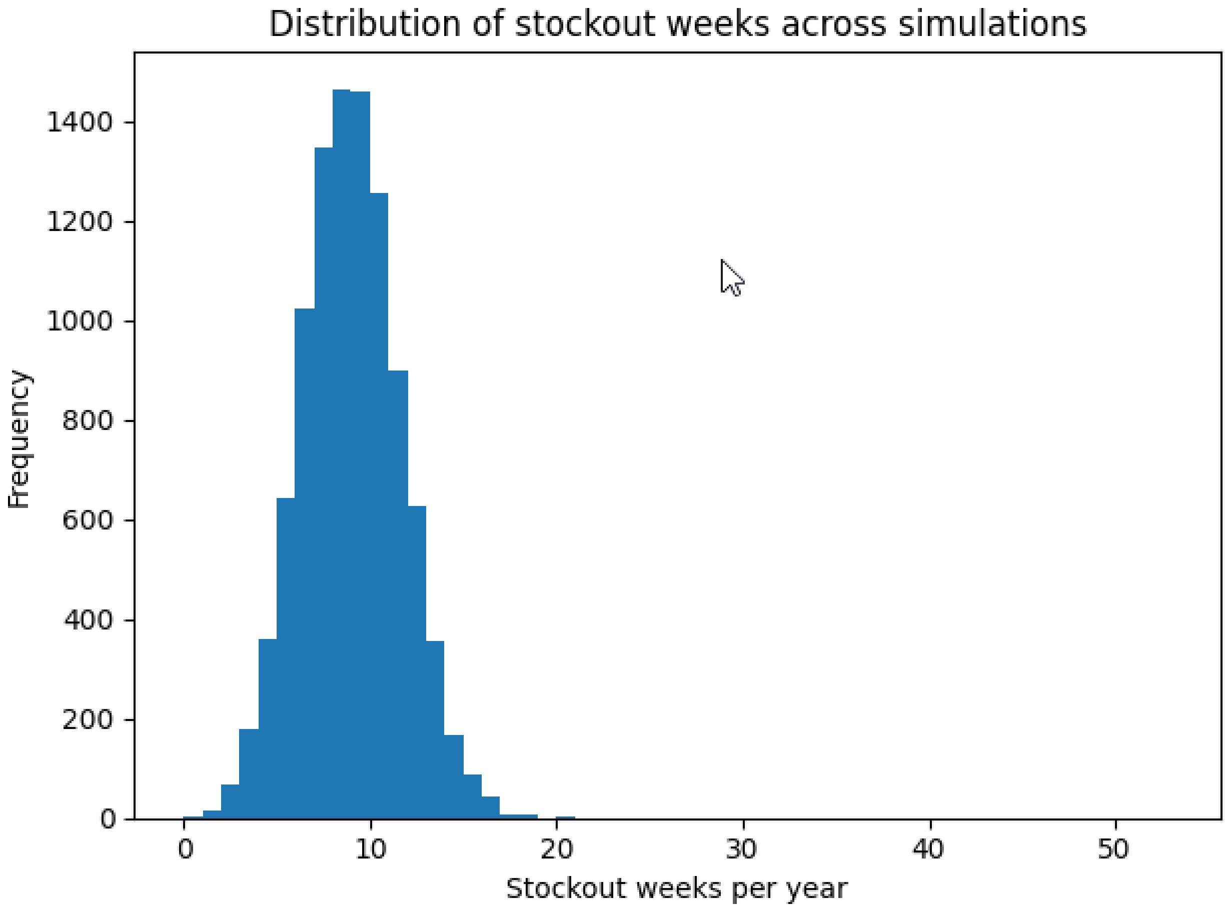

Figure 2: Distribution of stockout weeks per year (author's own simulation)

Figure 3 shows the distribution of final inventory levels at the end of the simulation period. The results indicate substantial variation in end-of-year inventory across scenarios. Some simulations end with low inventory levels, increasing the likelihood of stockouts, while others retain higher inventory buffers. This highlights a clear trade-off between reducing availability risk and increasing inventory holding costs. This trade-off suggests that improving availability through higher inventory levels reduces stockout risk but introduces additional cost and efficiency considerations.

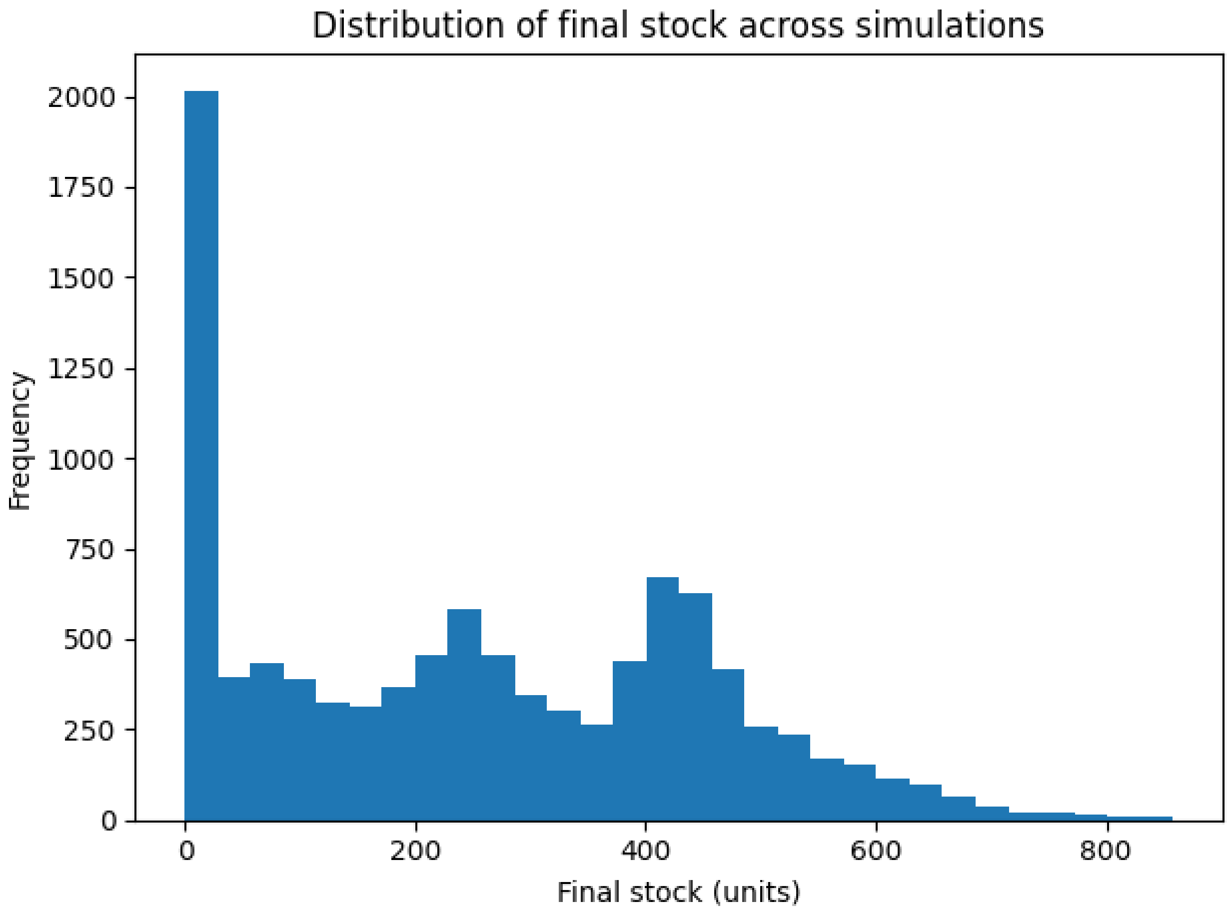

Figure 3: Distribution of end-of-year inventory levels (author's own simulation)

Figure 4 shows the distribution of defect probability across simulations. Most scenarios are associated with relatively low defect levels, while higher defect probabilities occur less frequently. This indicates that quality degradation affects a minority of cases rather than being a dominant risk. However, when quality degradation occurs, its impact is likely to be significant due to its potential effects on reputation and customer trust.

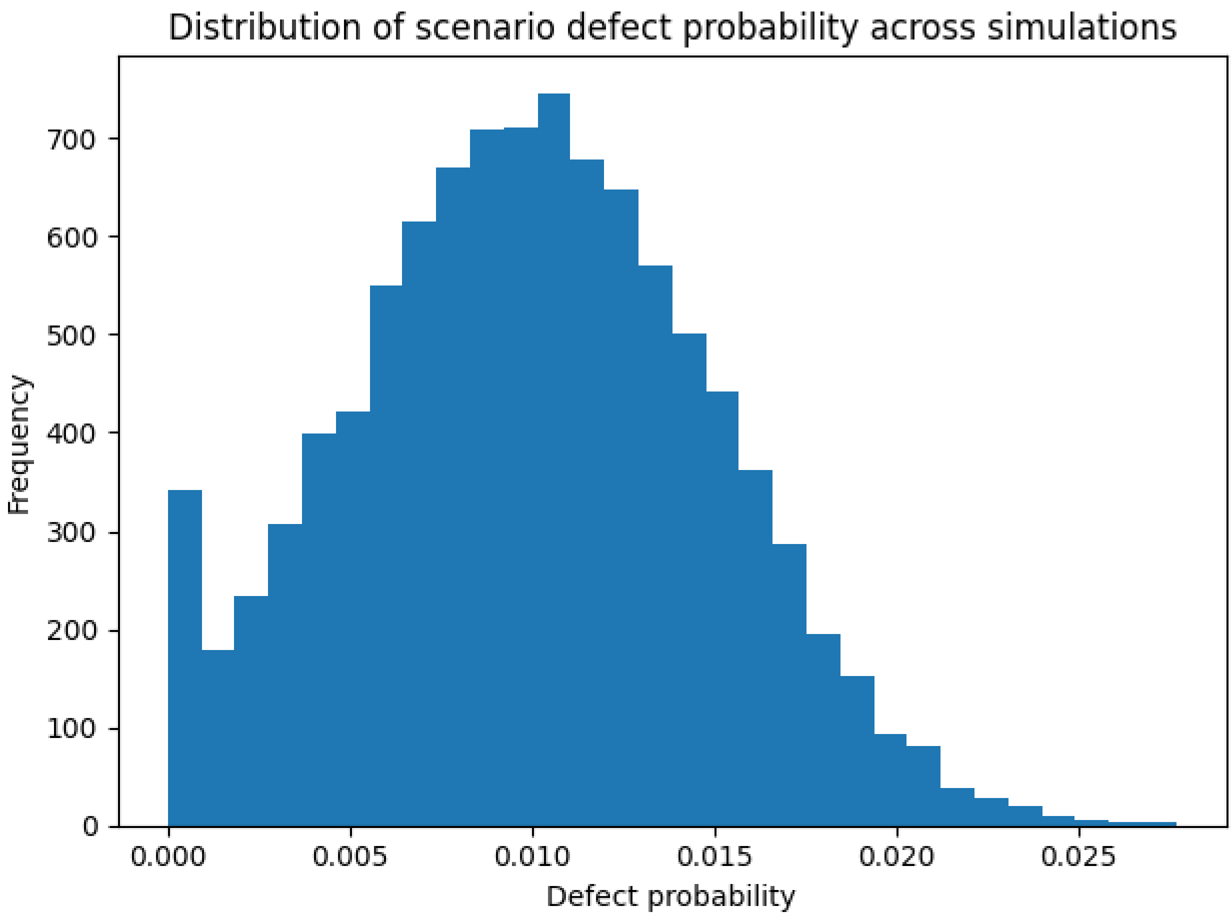

Figure 4: Distribution of quality defect probability (author’s own simulation)

Overall, the simulation results show that supply chain disruptions are likely to occur under uncertainty, primarily in the form of short-term availability issues, while quality degradation affects a smaller share of scenarios but represents a higher-impact risk. The results show that improving availability, therefore, requires accepting higher inventory levels and associated costs.

While the simulation indicates a very high probability of short-term stockout events over the annual period, this does not imply that the impact of such disruptions is acceptable for all customers. The probability estimates describe overall operational exposure under uncertainty, whereas the tolerance for availability failures is a governance decision. For high-profile customers, including the royal connection, product availability must always be maintained, even where this requires accepting higher inventory levels.

This represents an explicit management-level risk acceptance decision rather than an operational outcome of the quantitative model. Responsibility for availability failures affecting high-profile customers should therefore be understood as a management-level responsibility rather than a technical or operational one.

Business Continuity and Disaster Recovery Strategy

Business continuity planning should therefore focus on maintaining critical activities during short-term disruptions. Effective continuity depends on keeping critical services operational during disruption.

For Pampered Pets, critical services are those essential to sales, customer trust, and regulatory compliance. These include:

- Availability of the online shop, enabling customers to browse products and place orders
- Order processing, ensuring that customer orders can be handled and delivered

- Protection and integrity of customer and transaction data, which is essential for compliance with data protection requirements, including GDPR

Article 32(1)(b) and 32(1)(c) of the GDPR require organisations to ensure the ongoing availability and resilience of processing systems and the ability to restore access to personal data in a timely manner following a physical or technical incident.

To meet the stated business requirements, the disaster recovery strategy should support a recovery time objective (RTO) of less than one minute and a recovery point objective (RPO) of less than one minute.

As the business becomes more dependent on digital systems and external service providers, business continuity is closely linked to the choice of hosting platform. To support continuous availability and rapid recovery of critical services, Microsoft Azure is recommended as the hosting platform for Pampered Pets' digital operations. Cloud platforms with built-in geographic redundancy and automated recovery capabilities are widely recognised as effective enablers of business continuity and disaster recovery, particularly for customer-facing systems with strict availability requirements (Collier and Shahan, 2016). Azure provides native support for multi-region deployment, geographic redundancy, and automated failover mechanisms, making it suitable for meeting high-availability objectives (Microsoft, 2025a).

The proposed disaster recovery design is based on an active-active, multi-region architecture, in which customer-facing services operate concurrently across geographically separated regions. Microsoft Azure supports near-real-time data synchronisation through active geo-replication, enabling rapid recovery in the event of a regional outage (Microsoft, 2025b). Automated traffic management combined with continuous health monitoring enables traffic to be redirected to the available region within seconds. This supports a recovery time objective of less than one minute, while near-real-time replication limits potential data loss to seconds and ensures that the recovery point objective is met.

Reliance on a single cloud provider nevertheless introduces vendor lock-in risk, which may reduce organisational flexibility during disruption events. Such dependencies have been shown to increase response times during incidents and to complicate recovery when services cannot be migrated or replicated across alternative platforms at short notice (Ghasvari Jahrmoi et al., 2025). While this risk cannot be eliminated, it can be managed through architectural and contractual measures, including avoiding unnecessary dependence on provider-specific services, supporting data portability, and ensuring clearly defined exit arrangements. In this case, a controlled level of vendor lock-in is an acceptable trade-off to achieve the required levels of availability, resilience, and recovery performance.

From a disaster recovery perspective, the focus should be on maintaining minimum operational capability during disruption periods rather than on preventing all possible failures. This approach supports business continuity objectives while remaining proportionate to the organisation's size and maturity level.

While the proposed DR solution significantly reduces availability and data-loss risk, it does not fully mitigate dependence on cloud providers or eliminate risks related to cross-border outages and vendor-specific service failures. These residual risks require ongoing governance rather than purely technical controls.

Conclusion

Digitalisation and the move to an international supply chain increase the likelihood of disruptions to availability. In practice, this means short-term stockouts. They occur frequently under uncertainty but are usually brief.

Quality issues are different. They occur less frequently, but when they occur, the impact is greater.

With high-profile customers such as HRH the King and Prince Albert II of Monaco, both product quality and service availability must be treated as strategic requirements. Tolerance for failure is low. Even isolated quality problems or extended availability issues would have a disproportionate effect on reputation and customer trust.

The results point to a clear trade-off between improved availability, achieved through higher inventory levels, and increased costs. This reduces stockout risk but increases cost and reduces efficiency. Quality risk does not follow the same logic. It is best addressed through consistent process control across automated and geographically distributed operations, not through higher inventory levels.

Business continuity and disaster recovery, therefore, become central to the company's digital operations. Meeting the requirement for continuous service availability and minimal data loss requires RTO and RPO of less than one minute. Such recovery objectives involve higher costs but are considered proportionate to the reputational and availability risks being mitigated.

Digitalisation can support growth and efficiency for Pampered Pets, but only under clear conditions. Quality, availability, and continuity cannot be treated as optional. The focus should remain on resilience and control, not marginal efficiency gains.

References

Aijaz, M. and Nazir, M. (2024) ‘Modelling and analysis of social engineering threats using the attack tree and the Markov model’, International Journal of Information Technology, 16(2), pp. 1231–1238.

Aven, T. and Thekdi, S. (2025) Risk Science: An Introduction. 2nd edn. Oxon: Routledge.

Christopher, M. (2016) Logistics and Supply Chain Management. 6th edn. Harlow: Pearson Education Limited.

Calder, A. (2021) ISO 22301:2019 and business continuity management: Understand how to plan, implement and enhance a business continuity management system (BCMS). Ely: IT Governance Publishing.

Collier, M. and Shahan, R. (2016) Microsoft Azure Essentials: Fundamentals of Azure. 2nd edn. Redmond: Microsoft Press.

Demiray Kırmızı, S., Ceylan, Z. and Bulkan, S. (2024) ‘Enhancing inventory management through safety-stock strategies: A case study’, Systems, 12(7), 260. https://doi.org/10.3390/systems12070260

European Union (2016) Regulation (EU) 2016/679 (General Data Protection Regulation), Article 32: Security of processing. Available at: https://gdpr-info.eu/art-32-gdpr/ (Accessed: 8 January 2026).

Fadli, A.D. (2024) ‘Implementation of Total Quality Management (TQM) in improving product quality in manufacturing companies’, Maroon Journal of Management, 1(2), pp. 1–11. https://doi.org/10.37899/mjdm.v1i2.97

Fahimnia, B. et al. (2019) ‘Behavioral operations and supply chain management: A review and literature mapping’, Decision Sciences, 50(6), pp. 1127–1183.

Ghasvari Jahrmoi, M.A., Ashtiani, M. and Bakhshi, F. (2025) FaasFlows: an approach for reducing vendor lock-in and response time in serverless workflows. The Journal of Supercomputing, 81(1), Article 206.

Ivanov, D. and Dolgui, A. (2021) ‘A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0’, Production Planning & Control, 32(9), pp. 775–788. https://doi.org/10.1080/09537287.2020.1768450

Microsoft (2025a) Architecture strategies for disaster recovery. Microsoft Learn. Available at: https://learn.microsoft.com/azure/well-architected/reliability/disaster-recovery

(Accessed: 15 December 2025).

Microsoft (2025b) Business continuity and disaster recovery in Azure SQL Database. Microsoft Learn. Available at: https://learn.microsoft.com/azure/azure-sql/database/business-continuity-high-availability-disaster-recover-hadr-overview

Metropolis, N. (1987) ‘The beginning of the Monte Carlo method’, Los Alamos Science, Special Issue, pp. 125–130.

Sørensen, B.T. (2018) ‘Digitalisation: An Opportunity or a Risk?’, Journal of European Competition Law & Practice, 9(6), pp. 349–350.

Appendix A: Monte Carlo Simulation Code (Python)

```python
"""
Monte Carlo simulation supporting the Pampered Pets Executive Summary.

This script implements a Monte Carlo simulation to estimate the probability
of quality degradation and stockout events under uncertainty in demand,
lead times, and defect rates. All parameters and assumptions correspond
to those described in the Executive Summary and are documented inline
for transparency and reproducibility of the quantitative analysis.
"""

import numpy as np
import pandas as pd
import matplotlib.pyplot as plt

# ----------------------------
# Simulation parameters
# ----------------------------
NUM_SIM = 10000         # Number of Monte Carlo iterations
WEEKS = 52              # Simulation horizon (weeks per year)

# Weekly demand assumptions (normal distribution)
DEMAND_MEAN = 180       # Average weekly demand
DEMAND_STD = 30         # Demand standard deviation (COV ≈ 17%)

# Lead time assumptions (triangular distribution, weeks)
LEAD_TIME_MIN = 1
LEAD_TIME_MODE = 2
LEAD_TIME_MAX = 5

# Inventory policy parameters
INITIAL_STOCK = 800     # Initial inventory level
REORDER_POINT = 400     # Reorder trigger level
REORDER_QTY = 600       # Replenishment order quantity

# Quality assumptions
DEFECT_MEAN = 0.01          # Baseline defect probability
DEFECT_STD = 0.005          # Variability in defect probability
DEFECT_THRESHOLD = 0.015    # Threshold for quality degradation

# Random number generator (fixed seed for reproducibility)
SEED = 42
rng = np.random.default_rng(SEED)
```

Figure A1: Parameter definitions and assumptions used in the Monte Carlo simulation.

```python
# Monte Carlo simulation loop
# ----------------------------
for _ in range(NUM_SIM):
    # Draw scenario-specific defect probability
    defect_prob = rng.normal(DEFECT_MEAN, DEFECT_STD)
    defect_prob = float(np.clip(defect_prob, 0.0, 1.0))

    quality_degradation = defect_prob > DEFECT_THRESHOLD

    stock = INITIAL_STOCK
    outstanding_lt = None
    stockout_happened = False
    stockout_weeks = 0

    for week in range(WEEKS):
        # Receive replenishment if lead time has elapsed
        if outstanding_lt is not None:
            outstanding_lt -= 1
            if outstanding_lt <= 0:
                stock += REORDER_QTY
                outstanding_lt = None

        # Realised weekly demand
        demand = int(round(rng.normal(DEMAND_MEAN, DEMAND_STD)))
        demand = max(demand, 0)

        # Stockout occurs when demand exceeds available stock
        if demand > stock:
            stockout_happened = True
            stockout_weeks += 1
            stock = 0
        else:
            stock -= demand

        # Place replenishment order when reorder point is reached
        if stock <= REORDER_POINT and outstanding_lt is None:
            lt = rng.triangular(LEAD_TIME_MIN, LEAD_TIME_MODE, LEAD_TIME_MAX)
            lt = max(1, int(round(lt)))
            outstanding_lt = lt

    final_stocks.append(stock)
    scenario_defect_probs.append(defect_prob)
    stockout_weeks_list.append(stockout_weeks)

    quality_flags.append(quality_degradation)
    stockout_flags.append(stockout_happened)

    if quality_degradation:
        quality_degradation_count += 1
    if stockout_happened:
        stockout_count += 1
    if quality_degradation and stockout_happened:
        both_count += 1
```

Figure A2: Outcome combination matrix for stockout and quality degradation events.

```python
# Aggregate risk metrics
# ----------------------------
prob_quality_degradation = quality_degradation_count / NUM_SIM
prob_stockout = stockout_count / NUM_SIM
prob_both = both_count / NUM_SIM
prob_neither = 1 - (prob_quality_degradation + prob_stockout - prob_both)

avg_final_stock = float(np.mean(final_stocks))
avg_defect_prob = float(np.mean(scenario_defect_probs))
avg_stockout_weeks = float(np.mean(stockout_weeks_list))

# ----------------------------
# Reporting layer (tables + plots)
# ----------------------------
results = pd.DataFrame({
    "quality_degradation": np.asarray(quality_flags, dtype=bool),
    "stockout": np.asarray(stockout_flags, dtype=bool),
    "defect_probability": np.asarray(scenario_defect_probs, dtype=float),
    "final_stock_units": np.asarray(final_stocks, dtype=int),
    "stockout_weeks": np.asarray(stockout_weeks_list, dtype=int),
})
results["joint_quality_and_stockout"] = (
    results["quality_degradation"] & results["stockout"]
)

def fmt_pct(x: float) -> str:
    return f"{x:.2%}"

def fmt_num(x: float, decimals: int = 2) -> str:
    return f"{x:.{decimals}f}"

# Summary event table
event_table = pd.DataFrame({
    "Event": [
        "Quality degradation",
        "At least one stockout",
        "Joint quality and stockout",
        "Neither event",
    ],
    "Count": [
        int(results["quality_degradation"].sum()),
        int(results["stockout"].sum()),
        int(results["joint_quality_and_stockout"].sum()),
        int((~results["quality_degradation"] & ~results["stockout"]).sum()),
    ],
    "Probability": [
        results["quality_degradation"].mean(),
        results["stockout"].mean(),
        results["joint_quality_and_stockout"].mean(),
        ((~results["quality_degradation"] & ~results["stockout"]).mean()),
    ],
})
event_table["Probability"] = event_table["Probability"].map(fmt_pct)

```

Figure A3: Aggregation of simulation outputs and calculation of event probabilities.

```python

# Combination matrix table
combo_table = (
    results.assign(
        outcome=np.select(
            [
                results["joint_quality_and_stockout"],
                results["stockout"] & ~results["quality_degradation"],
                results["quality_degradation"] & ~results["stockout"],
            ],
            [
                "Both",
                "Stockout only",
                "Quality only",
            ],
            default="Neither",
        )
    )
    .groupby("outcome")
    .size()
    .rename("Count")
    .to_frame()
)
combo_table["Probability"] = (combo_table["Count"] / len(results)).map(fmt_pct)
combo_table = combo_table.reindex(
    ["Neither", "Stockout only", "Quality only", "Both"]
)

# Numeric summary table
numeric_table = pd.DataFrame({
    "Metric": [
        "Average final stock (units)",
        "Average defect probability",
        "Average stockout weeks per year",
        "Median final stock (units)",
        "95th percentile stockout weeks",
    ],
    "Value": [
        avg_final_stock,
        avg_defect_prob,
        avg_stockout_weeks,
        float(np.median(results["final_stock_units"])),
        float(np.percentile(results["stockout_weeks"], 95)),
    ]
})
numeric_table["Value"] = [
    fmt_num(numeric_table.loc[0, "Value"], 1),
    fmt_num(numeric_table.loc[1, "Value"], 4),
    fmt_num(numeric_table.loc[2, "Value"], 2),
    fmt_num(numeric_table.loc[3, "Value"], 1),
    fmt_num(numeric_table.loc[4, "Value"], 0),
]

```