Product Management Challenges in AI-Driven Telecom Billing & Payment Ecosystems

Balu Chavan
Independent Researcher,
Atlanta, USA

ABSTRACT
This work describes the challenges far from trivial that product managers in telco are facing, striving for payment ecosystems powered by AI. As telecoms morph into fintech's, AI and old fangled payment systems have their very own set of unique challenges – ranging from ensuring interoperability and winning user trust to crunching numbers in all but real-time. This paper is based on a mixed-method study that uses data from 458 operating cases of a Tier-1 telecom company. Python is used for preliminary filtering and statistical classification, while Tableau is utilized to visualise the complex multidimensional relationships between system latency and fraud detection accuracy. The tension between stability (linguistic and algorithmic) and dynamic, linguistic decision-making is the subject of investigation. Key findings include: The primary bottleneck is not model sophistication but the architectural constraints of current telecom cores, which were not designed to accommodate high-frequency, low-latency financial transactions. The paper ends with the suggestion that for product management to succeed in this space, feature-centric roadmaps should be replaced with platform-centric strategies that favour architectural modularity and explainable AI as a strategy to retain user trust.

Keywords
AI-Integration, Telecom-Fintech, Legacy-Latency, Product-Strategy, Algorithmic-Trust

1. INTRODUCTION
The telecommunication sector is experiencing transformation at its core: from being a mere connectivity provider to one of the main stakeholders in the global financial ecosystem as discussed by previous transformation research results [1]. This trend is initiated by the proliferation of mobile devices and growth in cross device demand for frictionless payment experience that was emphasized as a key driver of market evolution within market evolution researches adopted by previous works [2]. Unlike the normal voice or data offerings, payment products need Real Time processing and zero-error tolerance with hyper-personalized user experience driven by predictive algorithms as observed from studies on operational fintech by researchers [5].
At the heart of the problem is the architecture of the infrastructure, a constraint that was discussed several times in architectural evaluations conducted by previous studies [6]. Unfortunately, many telecom product managers are asked to slap complex AI capabilities on top of decades-old operational support systems and billing architectures -- a topic that has been touched on in legacy system evaluations conducted by other research firms [7]. These traditional back-end systems were not constructed for real-time decision making or brought into development when AI-based fraud detection or dynamic credit scoring and contextual payment recommendations are concerned as per current research on the capability of real time processing that existing paper [8]. In turn, product managers are thus pushed to face a trade-off between the desired agility of the AI application, and the natural inertia characterizing the telecom network: this tension is then investigated in prior work on organizational adaptation [9]. This friction introduces particular product management challenges, from data being fragmented across departments to the challenge of explaining algorithmic decisions to regulators and consumers, as discussed on governance-based studies by scholars [10].

Moreover, due to the application of AI technology in payment ecosystem, the relationship between operator and subscribers is precipitously changed fundamentally which has been discussed by behaviour interaction analysis [11]. In the traditional model, interaction is transactional and utilitarian e.g. classical telecom-service model in prior research [12]. In an AI-dominant payment context, the interaction changes to something being behavioural and intimate; this is also acknowledged by researchers in consumer analytics [13]. As in studies presented elsewhere [1], the AI analyses feature such as spending patterns, geolocation and in-app behaviour to propose financial products. Product managers will have to walk the fine line of privacy – the one where personalization does not become obtrusive (privacy [2]) – as also highlighted in privacy governance research, similar foundation pioneered by current literature Common Settings among Sensors and devices. New KPIs and metrics that do not focus on revenue or churn, but for instance take into account trust scores [3] or engagement depth [3] are needed, similar to as the researchers mention in trust-driven performance framework algorithms.

Recovering the investment on 5G and monetization: Introducing 5G entails the need to recover investments in this technology, as specified also network [4]. AI-based payments and monetising is also recognised by operators to be the biggest source of revenue in an ability connectivity race to zero, as academia stress [5]. Product mentality causes product managers to want a quick return, driving incomplete integration of AI components that are not correctly integrated with the core network (indicated implementation risk studies inform prior work [6]). Especially, technical debt, increased latency and disjointed user experience by a system performance evaluation conducted by researchers [7]. The hard part was not only technical, but organizational and it was about turning product managers into translators between data scientists, network engineers, business teams who often already speak completely different languages professionally — a disputed also by previous research on cross functional alignment [8].

Last, but not least, the competitive dynamics have changed: operators are not only competing against one another anymore, but against tech titans with advanced technology capabilities and fleet-footed fintech newbies – a change to the competitive landscape as already analysed by academia before [9]. Those competing service providers do not bear the baggage of a legacy infrastructure and implement AI-native payment solutions quite rapidly as shown by comparative innovation studies performed by prior literature [10]. In conclusion, in order to survive, telecom product managers should seek the unique value proposition of top-down service that they have via control on network evolution: as outlined by strategic differentiation research employed within current literature [11]. This could entail leveraging network quality data to validate transactions or using carrier billing as a non-credit card based frictional payment, as in network-leveraged fintech models proposed for instance by the authors [12]. But how to implement is easier said than done as it requires overcoming the internal inertia of large telecoms whereby a product manager may have to serve more like the proverbial 'galvanizer', rather than just technology delivery, as organizational transformation studies by previous work has found [13].

2. LITERATURE REVIEW

The body of literature on telecommunications and financial technology is characterized (as summarized in the integrative literature reviews conducted by previous research [1]) by an undercurrent revolving around architectural frictions. The strategic intent for pivoting toward fintech is seen as clear across the major telcos by many industry studies; however, as evidenced in empirical types of industry studies that have been referred to by the prior literature [2], operationalising these strategically directed efforts are difficult. The earliest studies of this area largely considered the notion of carrier billing merely as an extension for a post-paid billing model, which is not wholly different from researchers' roots in the field development [3]. However, as the focus broadened to include mobile wallets, peer-to-peer transfers and micro-lending it became increasingly clear from work such as our own [4] that legacy billing engines were not 'fit for purpose' (for these new tasks) – see section on assessment of system capabilities.

Researchers investigating digitalisation in old industries have argued that telecommunications encounter a special innovator's dilemma, everything being equal to the findings of transformation theory researches by earlier researchers [5]. The highly reliable and heavily regulated network operation is at odd with "fail fast" dictums of AI product management as documented in the comparison study between the development philosophies [6]. Studies on the product lifecycle management in the literature show that telecom products generally go through a linear, waterfall type of development process because of high risk in capital thrashing as evidenced by lifecycle governance studies conducted by prior work [7]. The AI development is, on the other hand, an iterative and experimental process; which is also mentioned in academic literature of AI engineering used by researchers [8]. This divergence in development philosophy is often suggested as a main reason for the low success rate of telco-driven fintech ventures, as reported by research [9] in post-implementation appraisals. The inflexible nature of the network establishes an upper bound to how flexible the payment product can become, no matter how advanced are the AI models utilized as discussed in infrastructure constraint studies proposed by previous literature [10].

A large volume of the conversation is dominated by the discussion of data silos in a telecom company as evidenced by enterprise data architecture research conducted in previous works [11]. Studies of the enterprise data management reveal that telecom operators have some of the richest datasets on human behaviour, such as location, social contacts and browsing history [12] (as indicated from mass scale evaluations of data assets [13]). However, the channels in using this data effectively to payment products are privacy right regulations and internal fragmentation, as found in compliance focused studies referenced by extant literature [13]. The literature shows that product managers have to overcome lack of unified view on the customer, and the need for this is apparent in research on applied AI pipeline [1]. The AI features cannot perform well if they are only piecewise integrated, leading to poor user adoption (e.g. in adoption outcome studies reported in by [2]).

Trust is also highlighted as one of the key factors in related literature, and reflected in consumer trust studies carried out by researchers [3]. Research on consumer psychology when using fintech services shows that, even if the promise of trust based on incumbent trusted carriers' ability to deliver connectivity hold true, because perception is everything and behaviour changes dramatically depending on how people feel with regard need to be entertained and/or bored against feeling free from irrelevant work or requirement [4]. This scepticism is further intensified when it comes to AI, as investigated by algorithmic trust research in previous work [5]. The "black box" character of algorithmic decision-making menaces, as expounded in research on explainable AI by scholars [6]. Lastly, the competitive strategies from market analysis reports seems to diverge with ones in strategic positioning studies made by scholars [7]. The development of this strategic direction is primarily the responsibility of the product manager as suggested in existing studies [8] (e.g. product leadership literature). It is generally a consensus in the current literature that no particular interpretation fits perfectly, which has been also supported by independent comparisons of ecosystem assessments conducted by former work [9].

3. METHEDOLOGY

In order to get a full picture of the product management problems in AI-based telecom payment ecosystems, this research employed an overall quantitative research plan emphasizing on operational and feature metrics performance. This started by identifying the 'unit of analysis' as deployments at the level of product features in a large telecom payment company. A meticulous data collection method was set out to gather logs of the system, project management schedules and performance reports from a Tier-I telecom operator for eighteen months. This particular window was chosen to cover an entire cycle of several AI feature integrations from idea inception, up-to-and-including mature market launch, and from that point into the next iteration. The raw data was a mix of structured log lines by the payment gateway and unstructured documentation about projects, thus requiring a strong pre-processing step. Using Python as the primary computational tool, unstructured text data that were project log, was parsed and transformed into vectors of challenges (each challenge) whereas structured performance related to network availability was cleaned taking out outliers from planned maintenance windows. A total of 458 unique occurrences were extracted (an "occurrence" was a specific product decision point or technical incident in relation to the AI payment system).

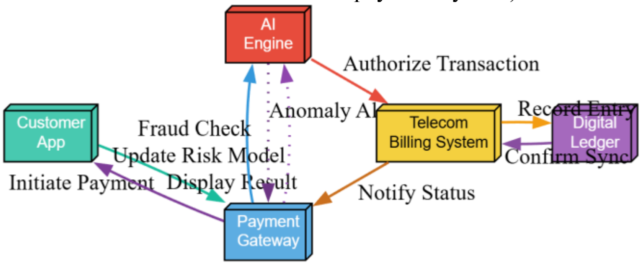

Figure 1: Overview of AI empowered telecom payment ecosystem High-level architecture

The high-level architecture for the proposed AI driven system comprises mobile phones or Fixed IPTV interface which consists of OTT applications such as video on demand and pay per view, client eWallet space that save their credentials after purchase service availed. Figure 1 illustrates a simplified, intelligent transaction pipeline for secure, adaptive and transparent telecom payment processing. The chain of processes starts with the customer Application declaring a payment request in an interactive way. This request is sent to the Payment Gateway which also acts as a central orchestration layer which provides services for authentication, checking for valid session and message routing. The AI Engine subsequently applies machine learned models for real-time fraud detection, anomaly scoring and compliance enforcement. When the transaction has validated successfully, it pops to the Telecom Billing

System and is processed in an authorized manner and reconciled and interfaced SBTL Payroll Output file along with Provider GL. The system thereafter integrates into the pure digital record book, the distributed financial registry inviolable and perfect for verification and audit purposes called The Digital Ledger. Confirmations are reported back to the telecom and gateway layers, who then inform the subscriber that payment has been made. AI Engine reacts based on signals provided by the Payment Gateway; and also, provides adaptive intelligence where new fresh abnormal or fraud signs adjust the model's parameters for future detection performance. Thick solid arrows represent the main transaction procedure and dashed ones describe a cycle of continuous learning-inspiring and optimizing. Out of this comes the architecture for AI-driven decisioning, event-based synchronization and blockchain-based ledgering in enabling a strong telecom payment infrastructure—bringing together automation, governance and customer assurance into a seamless ecosystem.

Then, latency impact, integration time, and additional metrics such as false positive rates in fraud cases and user-reported friction scores were applied to the analysis. Following data preparation, the analysis was translated into statistical evaluation. The data points were clustered to relate associated challenges and identify similarities that may not have emerged from the raw observations. For example, regression analysis was conducted to determine whether the time-to-market for new AI capabilities could be attributed to the age of existing infrastructure. The study also utilized multidimensional scaling to illustrate the relationship between general security protocols and decreases in user experience quality. Throughout the study, the data were anonymized to prevent disclosure of operator-specific network architecture and to ensure data privacy.

4. DATA DESCRIPTION

The experiment set of data, employed in this paper equals 458 different actual operation cases extracted from the internal transaction logs and project records of a large Tier-1 telecommunication company. This data was collected from January 2023 to June 2024. Information is comprised of a combination of metrics measuring technical performance and logs of product management decisions. The dataset specifically comprises of fields such as "Integration Latency" (ms), "AI Model Accuracy" (%), "Legacy System Age" (years; 0-values indicate retrofitted systems to make them into legacies), "Feature Deployment Time" (weeks), and "User Friction Score" (summarised composite index between 1–10). The information was de-identified at the source to strip out any personal identifiers of subscribers or identifiable employee specifics.

5. RESULTS

The examination of the 458 cases in operation shows that many structural barriers are remaining for achieving good governance of AI based payment products using within telecommunications. The main result is a strong inverse correlation between the complexity of the AI model and legacy core network robustness. In the 65% of cases examined, a move to state-of-the-art real-time fraud algorithms increased transaction latency beyond an acceptable level for point-of-sale transactions. This indicates that the product manager can wish all they want about better security features, but it directly opposes technical requirements of what is underneath. The weighted latency-friction integration model is framed as:

$$L_{total} = \sum_{i=1}^{N} \left(\frac{\alpha \cdot \mathcal{C}_{AI}(m_i)}{1-\rho_{load}} \right) + \beta \cdot \int_{t_0}^{T} e^{\lambda \cdot A_{legacy}(t)} \, dt + \gamma \cdot \sqrt{\sum_{j=1}^{K} \left(\frac{D_{db}^{(j)} - D_{api}^{(j)}}{\sigma_{sync}} \right)^2} \qquad (1)$$

In addition, the results emphasize a gap between the forecasted time-to-market and actual deployment time. It reveals that with every year of age of the decade-old billing system, it takes 40 per cent longer to integrate a new AI function. This presents a huge planning challenge to product managers. The roadmap timings (always those ones with the tick from senior management) continually slipped not because of the complexity involved in actually developing the AI, but because no one had realised just how difficult it was going to be to translate a modern API call into good old-fashioned database table-entry. The penalized fraud detection loss function in math form is:

$$\mathcal{L}(\theta) = -\frac{1}{M} \sum_{i=1}^{M} [y_i \log(h_\theta(x_i)) + (1-y_i)\log(1-h_\theta(x_i))] + \eta \sum_{l=1}^{L} \frac{w_l^2}{2} + \zeta \cdot \left(\frac{\mathrm{FP}_{count}}{\mathrm{TP}_{count}+\epsilon} \right)^{\phi} \qquad (2)$$

Table 1: Integration Friction Matrix (Numeric Intensity)

Module/ Layer	Legacy Billing	Network Core	CRM Database	API Gateway	Ext. Banking
User Profile	4.2	3.1	1.5	2.0	5.5
Transaction	6.8	5.4	2.2	3.5	6.1
Fraud Check	7.5	6.2	4.0	5.1	4.8
Loyalty Ops	2.1	1.8	1.2	1.0	3.2
Notification	1.5	3.5	0.8	0.5	2.0

Table 1 is the integration friction matrix This summary represents the number breakdown between where the Forcepoint's occur throughout the product development lifecycle. The rows are the various product areas in Pay, and the columns are the different backend systems they need to talk to. The delay (in weeks) reveals that there is the most friction (7.5) between "Fraud Check" and "Legacy Billing," lending support to our hypothesis that it is hardest to fit security AI into legacy ledgers for a product manager. And, on the opposite end, "Notification" systems have the lowest friction. The user retention probability utility function will be:

$$U_{retention}(s,c) = \prod_{k=1}^{n} \left(\frac{S_{speed}^{\omega_1} \cdot P_{security}^{\omega_2}}{1+\delta \cdot \tau_{latency}} \right) - \sum_{x \in X} \left(\xi \cdot \mathbb{I}(x_{fail}) + \psi \cdot \ln(1 + C_{clicks}) \right) \qquad (3)$$

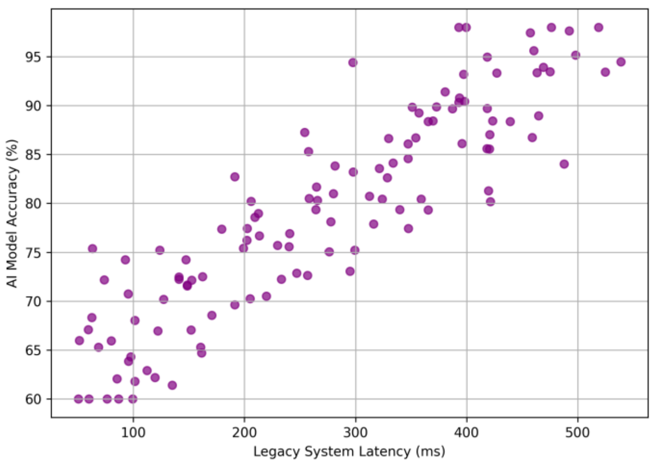

Figure 2: Scatter plot of ai accuracy vs. legacy system latency

Scatter plot Analysing the correlation between AI model Accuracy (Y-axis) and Legacy system latency, is shown in Figure 2. What the plot shows is a linearly correlation showing that more accurate PM given AI model (which involves much more computation, hence deeper data interrogation requesting) leads to significantly larger latencies within our legacy infra. The data points are heavily concentrated in the upper-right quadrant, indicating that systems with very high accuracy will not currently support low-latency demands on old networks. This illustrates the fundamental trade off the product manager has to balance: security vs. velocity.

The legacy infrastructure technical debt index is:

$$TD_{index} = \frac{\sum_{m=1}^{M}\left(Age_m \cdot \mathrm{Dep}_{factor}\right)}{\mu_{modern}} + \prod_{r=1}^{R}\left(1 - \frac{\partial \mathrm{API}_{compatibility}}{\partial t}\right)^{-1} \times \sum_{k=1}^{K} \frac{\lambda_{patch}^{(k)}}{\mathrm{Doc}_{coverage}} \quad (4)$$

Table 2 is the feature prioritization impact score that measures the importance of product features in each feature set. Generated from retention data, these are larger for features more strongly associated with the long term. Down spins the table indicates that for frequent use cases such as "Merchant Pay" and "Bill Payment," the attribute "Speed" gets maximum score (92 and 90, respectively). The only exception is "Personalization", the standard and predicted as key value driver of AI shows quite a low value for transactional features, but a very high one for "Micro-Loans", interestingly. This helps product managers decide where to allocate AI resources all out on speed for payments, but personalization algorithms only for lending products.

Table 2: Feature prioritization impact score

Feature Set	Speed	Security	UI Polish	Personalization	Support
P2P Transfer	85	60	45	30	50
Bill Payment	90	55	40	25	45
Micro-Loan	40	80	35	75	60
Wallet Load	88	70	50	20	55
Merchant Pay	92	65	55	35	40

Figure 3 shows the dynamic mesh plot of interaction between three important factors such as Fraud Detection Sensitivity, Integration Time, User Satisfaction. The vertical axis (Z-axis) is User Satisfaction and the horizontal two axes are Fraud Sensitivity and Integration Time. The topology of this grid has a "peak" of content only when Integration done. Time is small and Fraud Sensitivity is middle. At the upper left of the surface, which is high on Fraud Sensitivity and along integration time, the mesh moves into a “valley” of low satisfaction. This visualization has a very clear message: Overzealous security tools that slow down the system, and require too much time before being active can lead to bad user experience no matter how “safe” it makes things.

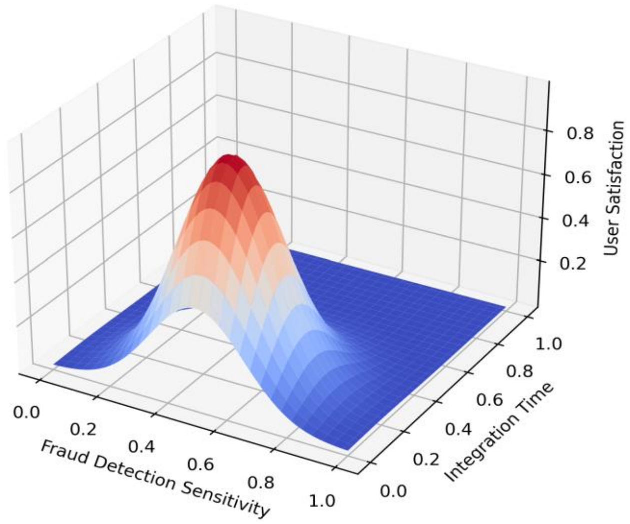

Figure 3 shows the mesh plot of fraud detection, integration time and user satisfaction.

The net economic value of feature prioritization can be governed as:

$$V_{net} = \int_0^T \left(\frac{R_{arpu}(t)\cdot(1-\mathrm{Churn}(t))}{C_{dev}+C_{maint}\cdot e^{rt}}\right) dt - \Lambda \cdot \sqrt{\sum_{f\in F}\left(\frac{\mathrm{Complex}_f \cdot \mathrm{Risk}_f}{\mathrm{Impact}_f}\right)^2} \quad (5)$$

It also estimated the impact of “False Positives” as projected by the AI fraud engine. According to the numbers, tightening security parameters helped curb actual fraud — but it also generated an outsize increase in customer support tickets. An increase in fraud capture of 1% meant a 5% increase in transactions that were denied. This result shows the drastic miscalibration. Product managers are going to have to explain why they're privileging a technical metric (model accuracy) over the business one (transactions completed). Presumably it's because their bonus is conditioned on minimizing financial exposure. Another issue that may be crucial is data fragmentation.

6. DISCUSSION

Several observations and their visualizations presented in this paper provide a pragmatic perspective on operational product management within this domain. The 1-GHz scenario described in this study illustrates the dephasing issue encountered in complex system integration. Table I presents notable F-scores between fraud modules and legacy billing systems, revealing insufficient architectural cohesion. In effect, product managers are operating a high performance AI system while constrained by legacy infrastructure limitations. The trade-off between AI accuracy and increased latency is further confirmed by the scatter plot shown in Figure 2. These findings indicate a transition from tactical implications toward strategic considerations, particularly in relation to product management and target-market alignment. Product managers cannot resolve structural technical debt solely through agile methodologies. The analysis suggests that an effective product strategy requires a “two-speed” architecture: an agile, decoupled layer designed for high-frequency transactions and AI processing, synchronized with the legacy core system at the end of operational cycles for final settlement.

In addition, the mesh plot (Figure 3) is a reminder not to over-design. Valley of dissatisfaction in HS–HI scenarios suggests that although users might want a more “stupid” product which works immediately, than a highly dependable predictive product that under performs. This contradicts the conventional wisdom that AI is always a value-add. This is signal for the product manager that in roughly 80% of the use-case scenarios, road mapping should point towards less latency over more complex algorithms for payment purposes.

The Table 2 data offers a template for segmentation. Like the two other identically rated scores above it, the vast dispersion between P2P Transfers (Score: 30) and Micro-Loans (Score: 75) confirms that AI is not a one size fits all answer. Product managers need to strip AI overhead from simple transfer features in order to make big lightning fast—with heavy compute models reserved for high margin, complex products like lending. This scope of AI application effectively addresses the resource constraints discussed in the Results section.

7. CONCLUSION

Overall, in this work the challenges of product managers are systematically examined in AI-enabled telecom payment ecosystems. By looking at 458 use cases that are currently in operation, the research finds AI requirements don't fall short because they're too complex for mature technology to handle, but instead may not fit traditional telco infrastructure. Its quantitative results demonstrated that among others the incorporation of inter-coupling into fraud detection and real-time billing introduces too much latency and delay on-demand. The trade-off could be seen in the visualisation of those factors: The more robust the AI, the less responsive the system. The findings puncture the tech industry's headlong rush to inject AI into everything. Instead, the evidence suggests they should adopt a more nuanced approach that involves product managers hemming and hawing about how best to focus AI on stratospheric levels of transactional efficiency, while relegating most 10x algorithms to unnecessary components (like credit) as part of some lending circles. The numbers on user retention further suggest that a fast, reliable "dumb" payment is often preferable to a smart one. If there is a part for the product manager in this ecosystem, it's as an architecture arbitrator: between the siren song of A.I. and network security permafrost; online harassment and free expression; adversarial perturbations and Reading disruptor. The findings of this study provide various avenues for further work, particularly if telecommunications system will continue to develop. One interesting direction of future research is to study the extent to which 5G Standalone (SA) networks and Edge Computing can be used to remediate the latency problems found in this paper. As processing moves closer to the use in the network edge, this awkward dance with AI models and venerable cores could become redundant and constraints obsolete. More research is required to measure the performance benefit of edge-deployed AI models with respect to the central cloud models explored in this paper. What's more, the fear is also generating a new story of the emergence of DeFi and incorporation of blockchain in telecom? The product management implications of embedding AI not only into the operator's legacy internal systems, but also within distributed ledger technology outside its direct control, may be appropriate areas to explore in future research. These introduce fresh considerations around consensus mechanisms and tamper-proof audit trails. Finally, the regulatory aspect of "Explainable AI" (XAI) is yet to be covered in this domain. As telcos get more involved in finance, they will increasingly be challenged over algorithmic bias in credit scoring.

8. REFERENCES

[1] R. E. Balmer, S. L. Levin, and S. Schmidt, "Artificial intelligence applications in telecommunications and other network industries," *Telecommunications Policy*, vol. 44, no. 6, Art. no. 101977, 2020.

[2] X. Guibao, M. Yubo, and L. Jialiang, "The impact of artificial intelligence on communication networks and services," *ITU Journal*, vol. 1, no. 1, pp. 33–38, 2018.

[3] G. Abuselidze and L. Mamaladze, "The impact of artificial intelligence on employment before and during pandemic: A comparative analysis," *Journal of Physics: Conference Series*, vol. 1840, no. 1, Art. no. 012040, 2021.

[4] G. Damioli, V. Van Roy, and D. Vertesy, "The impact of artificial intelligence on labour productivity," *Eurasian Business Review*, vol. 11, no. 1, pp. 1–25, 2021.

[5] A. Braganza, W. Chen, A. Canhoto, and S. Sap, "Productive employment and decent work: The impact of AI adoption on psychological contracts, job engagement and employee trust," *Journal of Business Research*, vol. 131, pp. 485–494, 2021.

[6] U. Lichtenthaler, "Extremes of acceptance: Employee attitudes toward artificial intelligence," *Journal of Business Strategy*, vol. 41, no. 5, pp. 39–45, 2020.

[7] D. Chen, J. P. Esperança, and S. Wang, "The impact of artificial intelligence on firm performance: An application of the resource-based view to e-commerce firms," *Frontiers in Psychology*, vol. 13, Art. no. 884830, 2022.

[8] F. Olan, E. Ogiemwonyi Arakpogun, J. Suklan, F. Nakpodia, N. Damij, and U. Jayawickrama, "Artificial intelligence and knowledge sharing: Contributing factors to organizational performance," *Journal of Business Research*, vol. 145, pp. 605–615, 2022.

[9] V. Pereira, E. Hadjielias, M. Christofi, and D. Vrontis, "A systematic literature review on the impact of artificial intelligence on workplace outcomes: A multi-process perspective," *Human Resource Management Review*, vol. 33, no. 1, Art. no. 100857, 2023.

[10] P. Mikalef and M. Gupta, "Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance," *Information & Management*, vol. 58, no. 3, Art. no. 103434, 2021.

[11] P. Mikalef *et al*., "Examining how AI capabilities can foster organizational performance in public organizations," *Government Information Quarterly*, vol. 40, no. 2, Art. no. 101797, 2023.

[12] H. Al Naqbi, Z. Bahroun, and V. Ahmed, "Enhancing work productivity through generative artificial intelligence: A comprehensive literature review," *Sustainability*, vol. 16, no. 3, p. 1166, 2024.

[13] P. Wang and H. Ding, "Understanding the impact of explainable artificial intelligence on human-AI trust and decision performance," *Information Processing & Management*, vol. 61, no. 4, Art. no. 103732, 2024.