Journal of Management and Planning Research
2024, 1(1), 139-148

STRATEGIC CHANNEL DESIGN AND TRANSACTION-COST PLANNING IN THE LIGHTNING NETWORK

Alexander Tzonis

The Lightning Network has become a leading operational response to Bitcoin's throughput and fee constraints. For managers and planners, the central question is a design choice: when should two parties continue to settle directly on-chain, when should they rely on a unidirectional channel, and when is a bidirectional channel the economically superior settlement architecture? This paper develops a management-facing analytical guide to that decision. Using the formal results reported by Guasoni, Huberman, and Shikhelman, the analysis clarifies the trade-offs among transaction fees, locked collateral, payment frequency, and payment symmetry, and it reorganizes the underlying propositions into implementable planning rules. The paper shows that direct on-chain settlement remains efficient only when transaction intensity is very low; unidirectional channels become attractive as one-sided flow increases; and symmetric or near-symmetric bilateral traffic makes bidirectional channels substantially more efficient because netting reduces both long-run cost and blockchain congestion. The discussion highlights concrete parameter-driven implications using the published benchmark settings (including channel-reset costs proportional to transaction size and a one percent interest rate) and provides a practical decision framework that can be applied to budgeting, liquidity commitment, and capacity planning. The resulting contribution is an operationally interpretable set of channel-design heuristics for payment-system managers, infrastructure planners, and researchers concerned with digital transaction design, cost-efficient settlement, and throughput management.

https://doi.org/10.66033/jmpr2024-112

INTRODUCTION

Management and planning research is centrally concerned with the design of systems that allocate scarce resources efficiently under recurring operational constraints. In digital payments, one such constraint is settlement capacity: Bitcoin processes base-layer transactions on-chain, but protocol-level throughput limits imply that settlement can become expensive and slow when demand is high [10, 7]. The Lightning Network (LN) mitigates this bottleneck by enabling counterparties to execute frequent transfers off-chain while preserving cryptographic security and enforceability [12, 11].

For organizations that accept, route, or batch digital payments, the resulting design problem is managerial rather than purely technical. It requires deciding when it is economical to commit capital, how much collateral should be locked to maintain service levels, and which settlement architecture minimizes long-run cost for a given flow profile. The economic structure of this choice is analyzed rigorously by [5], who derive explicit cost functions, characterize threshold conditions for channel use, quantify optimal collateral choices, and connect channel design to blockchain congestion.

This manuscript contributes by translating those formal results into a planning-oriented decision framework that is directly usable in managerial contexts. Specifically, it (i) restates the model primitives in operational terms (fees, intensity, asymmetry, and collateral cost), (ii) consolidates the underlying propositions into interpretable decision regions that separate when on-chain settlement, unidirectional channels, or bidirectional channels are cost-minimizing, and (iii) provides an implementation-ready guide for collateral and throughput planning that follows from the published benchmark illustrations. Throughout, the objective is clarity and actionability: the paper emphasizes how parameters map into decisions, what boundary conditions matter, and which simplifying assumptions are essential for interpreting the rules of thumb.

Consistent with this objective, the analysis is explicitly source-based: it relies on the benchmark settings, propositions, and numerical examples reported in the published study, and it validates the managerial rules by cross-checking that the qualitative decision regions are consistent across the analytical expressions and the reported numerical benchmarks. The contribution is therefore not a new theory of payment channels, but a rigorous operationalization of an established analytical result set for management and planning audiences, with careful attention to decision relevance, interpretability, and applicability.

LITERATURE CONTEXT AND OPERATIONAL BACKGROUND

The LN belongs to a broader family of arrangements in which frequent counterparties economize on transaction costs by settling periodically rather than after every exchange. This logic is closely related to classic inventory-theoretic and cash-management models developed by [1], [14], and [9]. In those traditions, the central question is how to balance transaction costs against the opportunity cost of holding resources idle. In the LN, the corresponding idle resource is collateral locked inside the channel.

The LN itself is a network of channels connecting pairs of nodes. A channel opens when both parties fund a jointly controlled account on-chain. Thereafter, payments are represented as off-chain balance updates, while the total funds committed to the channel remain fixed until closure and reopening [12, 15]. This design improves throughput by shifting repeated payment activity away from the blockchain.

The source study also identifies contemporary application settings that make the framework operationally meaningful. These include everyday retail-like flows, cross-border remittances, and digital gaming payouts, all of which depend on repeated transactions, cost sensitivity, and settlement speed [8, 13, 4]. In this sense, LN channel choice is an infrastructure planning problem: it requires matching system architecture to the volume, direction, and persistence of expected payment demand.

ANALYTICAL FRAMEWORK

Decision environment

Consider two nodes exchanging payments in both directions. Payments from node 1 to node 2 arrive according to a Poisson process with rate λ_1, and payments from node 2 to node 1 arrive according to an independent Poisson process with rate λ_2 [5]. The continuously compounded interest rate is denoted by r, the cost of resetting a channel by B, and the cost of a direct on-chain transaction by C. These primitives map cleanly into managerial levers: λ_1 and λ_2 summarize demand intensity and directional imbalance; r captures the opportunity cost of tying up working capital; and B and C represent the two fee-like frictions that determine whether repeated transactions should be pooled off-chain or settled individually on-chain.

The source analysis assumes unit transaction size for tractability. As explained by [5], this can be interpreted as a reduced-form representation of repeated payments of roughly comparable size, or equivalently as working in transaction units where the mean payment size is normalized. The Poisson structure is likewise a deliberate simplification: it provides a transparent way to separate intensity (how often payments arrive) from asymmetry (how unbalanced the two directions are), which are the two most decision-relevant drivers of whether netting inside a channel creates savings.

Settlement architectures

The relevant planning options are:

1. **All-on-chain settlement**, in which every payment is posted directly to the blockchain.
2. **A unidirectional or hybrid arrangement**, in which the more frequent payer uses a channel while the reverse direction may still rely on on-chain settlement when the channel cannot support it.
3. **A bidirectional channel**, in which both parties commit collateral and use a single shared channel.

The baseline on-chain cost is explicit and linear in transaction intensity:

$$J_{\text{on}}(\lambda) = \frac{C\lambda}{r}. \tag{1}$$

This expression provides the benchmark against which channel-based alternatives are evaluated [5].

What the analytical model solves

The principal managerial quantities identified in the source study are:

- exact channel cost as a function of transaction rates and chosen collateral,
- cost-minimizing collateral levels,
- threshold conditions under which channels become economical,

- the long-run ratio of off-chain to on-chain transactions, and
- the probability that one side exhausts its balance before the other.

These outputs translate into concrete planning questions: how much liquidity should be committed to sustain a target level of off-chain throughput, when should a firm pay the fixed reset cost rather than continuing to incur per-transaction fees, and how does directional imbalance change the risk of premature exhaustion and therefore the frequency of resets. In this sense, the model provides an analytically grounded bridge between transaction-flow forecasting and operational decisions about settlement architecture, collateral budgeting, and expected blockchain load.

SOURCE-BASED ANALYTICAL RESULTS FOR CHANNEL PLANNING

When channels are economically justified

The general qualitative result is clear: channels become attractive when on-chain costs are high, channel-reset costs are low, payment rates are sufficiently high, and interest rates are low [5]. This is the core threshold logic of the model and is the first-order planning rule for system designers.

Equally important, the source paper shows that the benefit of channelization is strongest when payment frequencies in both directions are equal or nearly equal. When flows are materially imbalanced, the economic character of the arrangement becomes closer to a unidirectional channel than to a truly symmetric bidirectional one [5].

Asymptotic cost structure

For operational planning, the source paper's asymptotic results are especially useful because they summarize how cost and required collateral scale in the economically relevant case of low interest rates. Table 1 restates the key results in a form designed for managerial interpretation.

Table 1: Source-based asymptotic implications for settlement design

Configuration	Applicable payment pattern	Cost / capital order	Planning implication
All-on-chain	Very low transaction intensity	Linear in payment volume	Best when counterparties transact too rarely to justify locked capital.
Unidirectional channel	Strongly one-sided payment flow	Minimal cost and required capital scale with $r^{-1/2}$	Useful when one party pays much more frequently than the other.
Symmetric bidirectional channel	Equal bilateral payment rates	Minimal cost and required collateral scale with $r^{-1/3}$	Most efficient architecture because netting lowers effective settlement pressure.
Asymmetric bidirectional channel	Both directions active, but materially unequal	Operationally behaves like a unidirectional design at leading order	A shared channel may still be used, but its economics are dominated by net flow rather than gross two-way volume.

Notes: The source study shows that the minimal cost of a unidirectional channel grows with the square root of payment rates, while the minimal cost of a symmetric bidirectional channel grows with the cubic root of payment rates. This is the analytical reason that payment symmetry materially improves efficiency [5].

For the two most important special cases, the source paper provides explicit asymptotic expressions:

$$c_{\text{opt}}(0,\lambda) \approx 2\left(\frac{B\lambda}{r}\right)^{1/2} - \frac{B}{2}, \tag{2}$$

$$c_{\text{opt}}(\lambda,\lambda) \approx 3\left(\frac{2B\lambda}{r}\right)^{1/3} - \frac{B}{6}. \tag{3}$$

The corresponding optimal collateral scales as

$$l_{\text{uni}} \approx \left(\frac{B\lambda}{r}\right)^{1/2}, \tag{4}$$

$$l_1 = l_2 \approx \left(\frac{2B\lambda}{r}\right)^{1/3}. \tag{5}$$

The operational meaning is straightforward: once the environment supports channel use, symmetric bilateral traffic becomes relatively more efficient than one-sided traffic because cost and capital increase more slowly.

Benchmark decision regions from the published numerical illustration

The published study includes a benchmark numerical illustration (reported as Figure 1 in [5]) that is particularly useful for decision-making because it translates the analytical expressions into a visible partition of the (λ_1,λ_2) space. In that benchmark, the horizontal and vertical axes represent payment rates λ_1 and λ_2 over the range 0 to 0.2, the round-trip channel-reset cost is normalized to $B = 1$, and the interest rate is set to $r = 1\%$. The managerial value of this benchmark is twofold: it provides an immediate "sanity check" that the qualitative comparative statics implied by the model are reflected in the reported numerical boundaries, and it offers a simple operational map from observed flow patterns to the least-cost settlement architecture.

Under the benchmark, multiple configurations are feasible in principle (including on-chain settlement, mirror-image hybrid cases, separate unidirectional channels, and a single bidirectional channel). However, the reported comparison across all feasible configurations indicates that two separate unidirectional channels do not emerge as the least-cost design once all costs are accounted for [5]. The resulting practical decision regions are summarized in Table 2.

Table 2: Operational decision guide under the benchmark illustration reported by [5]

Observed payment pattern	**Least-cost architecture**	**Managerial meaning**
Both payment rates low	All-on-chain settlement	Counterparties that transact only very rarely should avoid locking capital.
One rate high, the other low	Unidirectional / hybrid channel	The more frequent payer benefits from channel use; the less frequent side can continue to rely on on-chain settlement when needed.
Both rates sufficiently high	Single bidirectional channel	Shared collateral and netting make a bidirectional design the dominant planning choice.
Near-symmetric bilateral traffic	Bidirectional channel with strongest efficiency gain	Netting materially lowers effective cost, and the design becomes relatively insensitive to moderate increases in the lower rate.

Notes: In the published benchmark, the study emphasizes that, when $B = 1$ and $r = 1\%$, a symmetric bidirectional channel is already economical at very modest frequencies, including the case $\lambda = 0.2$, interpreted by the authors as roughly one transaction every five years in each direction [5].

Interpreted operationally, the benchmark confirms three robust planning messages. First, at low overall intensity, the fixed reset cost and financing cost dominate and all-on-chain settlement is preferred. Second, as intensity rises in a strongly imbalanced corridor, a unidirectional or hybrid design can be cost-minimizing because the dominant payer gains most of the amortization benefit. Third, as bilateral traffic becomes sufficiently frequent and near-symmetric, a single bidirectional channel becomes the dominant planning choice because netting reduces both the required collateral buffer and expected reset frequency. This benchmark is therefore valuable not because it introduces new assumptions, but because it transparently validates the architecture-switching logic and illustrates that, under favorable cost-of-capital conditions, channel use can be justified at much lower frequencies than casual intuition would suggest.

A concrete collateral-planning example

A central operational insight from the framework is that collateral is not merely a technical requirement; it is a planning variable that trades off service continuity against the opportunity cost of locked capital. The numerical example reported in the source study is therefore best interpreted as a collateral-budgeting exercise: given an expected payment intensity and a targeted architecture, the planner chooses collateral to minimize the sum of reset costs and financing costs [5].

From a managerial standpoint, the key takeaway is the shape of the cost function. Too little collateral increases the frequency of resets (and therefore the expected number of on-chain interactions needed to reopen channels), while too much collateral increases financing costs without proportionate reliability gains. The optimal point balances these forces, and the example demonstrates that even modest transaction rates can justify meaningful collateral commitment when on-chain fees are nontrivial and when bilateral netting is available. This logic also suggests a straightforward operational procedure: forecast (λ_1, λ_2) over the relevant planning horizon, select the feasible architecture, and set collateral based on the cost-minimizing level implied by the model rather than by ad hoc safety margins.

Congestion reduction and throughput planning

Beyond private cost minimization, the framework clarifies how channel adoption changes system-wide load. A channel substitutes many off-chain balance updates for fewer on-chain interactions (open/close/reset), so the ratio of off-chain payments to on-chain transactions becomes an operational measure of throughput relief [5]. For planners, this metric has two implications. First, channel-rich architectures can be understood as capacity-expanding investments: they reduce the marginal use of the scarce blockchain "resource" per unit of economic activity. Second, the magnitude of congestion relief is not uniform; it depends on intensity and symmetry. When flows are symmetric, netting reduces the need for large unilateral collateral buffers, which further reduces resets and therefore further lowers on-chain load. When flows are highly asymmetric, the probability of exhausting one side increases, which partially offsets congestion gains by requiring more frequent resets.

These distinctions matter for operational policy. Organizations that route payments (e.g., platforms or exchanges) can prioritize bidirectional channels on high-volume, symmetric corridors to maximize congestion relief, while treating low-volume or highly imbalanced corridors as candidates for on-chain settlement or limited-capacity unidirectional channels.

A PRACTICAL DECISION FRAMEWORK

For management and planning audiences, the core findings can be represented as a simple decision process. Figure 1 translates the source paper's results into a strategic design guide.

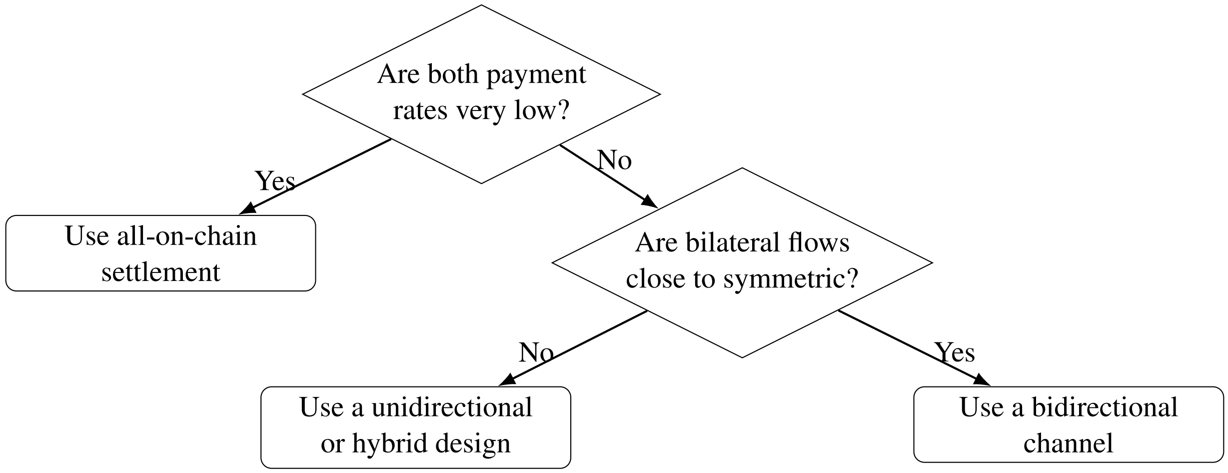

Figure 1: Management-oriented decision logic implied by the analytical results. Low transaction intensity favors direct settlement, asymmetry favors unidirectional or hybrid designs, and near-symmetry favors bidirectional channels.

This decision logic is intentionally simple, but it captures the source study's most durable conclusions and converts them into operational guidance.

IMPLICATIONS FOR PRACTICE AND ADVANCEMENT OF RESEARCH

Implications for managers and system planners

Three practical implications follow.

First, **payment frequency should be treated as a planning variable, not merely a descriptive one**. The value of a channel depends on the expected persistence of bilateral flow. Managers designing digital payment relationships should therefore distinguish between sporadic interactions and durable recurring flows.

Second, **payment symmetry is a strategic design criterion**. The source study shows that the economic gain from channelization is highest when bilateral flows are equal or nearly equal. In applied settings, this means that counterparties with balanced two-way exchanges are the strongest candidates for bidirectional infrastructure.

Third, **capital allocation should be planned asymmetrically when flow is asymmetric**. The monthly-versus-weekly example demonstrates that the heavier payer may need to commit several times more capital than the lighter payer. Practical channel design must therefore account for who carries the liquidity burden, not only whether the channel lowers total cost.

Competitive implications under high fee environments

The source paper also offers a concrete market-comparison example. When on-chain fees are low, the fee for a typical LN transaction can be as little as approximately $0.05. But if opening and closing a channel were to cost $50 each, implying a round-trip cost of $100, the channel would need to process roughly $3,334 to outperform a payment network charging a 3% interchange fee [5].

This calculation is especially relevant for management and planning research because it links internal system design to external competitive positioning. It shows that the competitiveness of LN-based settlement depends not only on protocol efficiency but also on broader fee conditions and transaction volume thresholds.

Implications for future research

The managerial framework presented here also highlights several research directions that would strengthen external validity without changing the core intuition. First, empirical measurement of (λ_1, λ_2) in real payment networks would allow researchers to map observed flow distributions into the decision regions and test whether organizations' channel choices align with the implied cost minima. Second, extending the model to heterogeneous transaction sizes would help connect the unit-size abstraction to environments with heavy-tailed payment distributions, where occasional large transfers can dominate collateral needs. Third, network-level extensions that incorporate multi-hop routing, competition among routing nodes, and endogenous fee setting would allow researchers to study how private incentives interact with social congestion outcomes.

Finally, future work could integrate operational constraints that managers routinely face but that are abstracted from in the benchmark model, such as risk limits on locked liquidity, regulatory capital constraints, and service-level requirements for payment latency. These extensions are not required to use the current framework as a planning heuristic, but they would help quantify how robust the decision regions are when real-world frictions and institutional constraints are layered onto the idealized settlement environment.

CONCLUSION

The Lightning Network is not only a cryptographic innovation; it is also a planning mechanism for allocating settlement capacity under recurring fee and liquidity constraints. The core managerial decision is whether repeated payment flows justify shifting activity off-chain and, conditional on that choice, which channel architecture minimizes long-run cost while maintaining operational reliability.

The analytical results organized in this manuscript support several decision-relevant conclusions. Direct on-chain settlement is most sensible when expected transaction frequency is very low, because fixed channel reset costs and financing costs cannot be amortized. Unidirectional or hybrid structures become attractive when flow is materially one-sided, because the dominant payer can internalize most of the benefit from reduced on-chain settlement without requiring symmetric collateral contributions. Bidirectional channels become the dominant design when both parties transact frequently and especially when traffic is symmetric or nearly symmetric; in that region, netting reduces both total expected cost and blockchain congestion by lowering the required collateral buffer and reducing the expected number of resets.

For management and planning audiences, the practical value of these results is interpretability: transaction intensity and symmetry are sufficient statistics that map directly into architecture choice, and collateral is a budgetable decision variable rather than a fixed technical parameter. At the same time, the framework should be applied with awareness of its simplifying assumptions, including unit transaction size and stylized arrival processes. These assumptions are valuable for producing transparent decision regions, but they also motivate the empirical and network-level extensions outlined above. Within those bounds, the paper provides a professional, implementable guide to transaction-cost planning in LN-style settlement systems and a clear baseline for future research that connects payment-network engineering to managerial capacity allocation.

REFERENCES

[1] Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. *Quarterly Journal of Economics*, 66(4), 545–556.

[2] Brânzei, S., Segal-Halevi, E., & Zohar, A. (2022). How to charge lightning: The economics of bitcoin transaction channels. In *2022 58th Annual Allerton Conference on Communication, Control, and Computing* (pp. 1–8).

[3] Frost, J., Shin, H. S., & Wierts, P. (2020). *An Early Stablecoin? The Bank of Amsterdam and the Governance of Money*. Bank for International Settlements Working Paper No. 902.

[4] Graf von Luckner, C., Reinhart, C. M., & Rogoff, K. (2023). Decrypting new age international capital flows. *Journal of Monetary Economics*.

[5] Guasoni, P., Huberman, G., & Shikhelman, C. (2024). Lightning network economics: Channels. *Management Science*, 70(6), 3827–3840.

[6] Huberman, G., Leshno, J. D., & Moallemi, C. (2019). An economist's perspective on the bitcoin payment system. *AEA Papers and Proceedings*, 109, 93–96.

[7] Huberman, G., Leshno, J. D., & Moallemi, C. (2021). Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *Review of Economic Studies*, 88(6), 3011–3040.

[8] Ibaba, A., Abdulmalik, A., Alhassan, A. B., Sunusi, I., & Thomas, S. (2021). Rem9ja: Bitcoin lightning network remittance solution. In *2021 1st International Conference on Multidisciplinary Engineering and Applied Science* (pp. 1–4).

[9] Miller, M. H., & Orr, D. (1966). A model of the demand for money by firms. *Quarterly Journal of Economics*, 80(3), 413–435.

[10] Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*.

[11] Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press.

[12] Poon, J., & Dryja, T. (2015). *The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments*.

[13] Tetek, J. (2021). Strike: Disrupting the remittance payments via bitcoin's lightning network.

[14] Tobin, J. (1956). The interest-elasticity of transactions demand for cash. *Review of Economics and Statistics*, 38(3), 241–247.

[15] van Wirdum, A. (2016). Understanding the lightning network.

Alexander Tzonis, Department of Architecture, University of Technology Delft, 108 Oude, Delft 2611, Netherlands

Manuscript Published; 22 November 2024.