Stochastic Routing for Cross-Border Payments

Marco Paes
Computer Science Major (Junior)
marcorp@stanford.edu

Andrea Nam
Computer Science Coterm
anamsong@stanford.edu

***Abstract*—This paper presents a greedy stochastic model for identifying routes in cross-border money transfers over a heterogeneous network of banks. We study routing under uncertainty in a graph where each node indicates which bank currently holds the funds and in what currency. Edges correspond either to intra-bank FX conversions or inter-bank transfers in the same currency. Exchange rates evolve stochastically around commercial benchmarks, and each institution applies its own spreads and fees. Given a user directive specifying a sender bank, source currency, receiver bank, target currency, and amount, the model constructs feasible multi-hop routes and evaluates actions using a Monte Carlo one-step lookahead policy. At each decision point, the agent samples institution-specific markups and fees, filters out actions that cannot reach the destination, and greedily selects the move with the highest expected value in the destination currency, adjusted for progress in the network. Over many simulated FX scenarios using a decade of historical data, we obtain a distribution of realized transfers and benchmark them via a fairness ratio comparing the final amount to a direct commercial conversion. Our results show that, under realistic spreads and corridors, institution-aware multi-hop routing can systematically outperform direct conversions in both cost efficiency and value delivered to the recipient.**

I. Introduction

Financial transfer networks are dynamic and stochastic, and the choice of how and where to route a transfer can significantly affect both cost and reliability. Motivated by our own experiences as international students from Latin America encountering slow, expensive, and opaque remittance routes, this project examines cross-border payment routing over a network of banks and currencies, where each bank applies its own exchange-rate markups, fee schedules, and corridor-specific constraints. The objective is to identify a path that maximizes the value ultimately delivered to the recipient by navigating institutional constraints, FX volatility, and the transaction costs incurred along the way.

In practice, a cross-border transfer is not just a single FX operation. The sender initiates a payment from a specific bank and currency (e.g., SFCU in USD), funds travel through a sequence of correspondent institutions via same-currency transfers, and at one or more points some institution converts from the source to the destination currency. Each institution maintains accounts only in a subset of currencies and supports a limited set of internal FX pairs; some banks cannot hold or convert certain currencies at all (e.g., many US and European institutions do not support PYG-denominated accounts). Hence, agent must track both which bank currently holds the funds and in which currency they are held. Routing decisions therefore operate over this joint institution–currency state, which determines which transfers or FX actions are feasible at each step

We study this problem using an explicit institution-currency graph. Nodes represent $(\text{institution}, \text{currency})$ pairs, and edges come in two types:

1) Transfer edges: Funds move between banks in the same currency and incur a flat transfer fee
2) FX edges: An institution converts between two of its supported currencies at a stochastic effective rate derived from commercial benchmarks plus a bank-specific markup and percentage fee.

Uncertainty arises from the evolution of commercial FX rates and from random bank-to-market spreads around those benchmarks.

In the following, we formalize the problem over an institution-currency graph, describe the data and stochastic environment, define the greedy routing policy and fairness metric, report empirical results, and then compare this approach conceptually to the naive and POMDP-based alternatives we initially explored. We also suggest directions for extending this framework toward richer stochastic control formulations in financial networks.

II. Problem Statement

Cross-border money transfers typically occur through a network of correspondent banks, each applying its own exchange rates, margins, and transaction fees. A conventional transfer from a sender in currency s to a recipient in currency t is computed as a single logical route chosen by the bank or payment provider. This route may traverse multiple intermediaries, but the user usually has no control over which institutions are involved, how many times the funds are converted, or how costs are allocated along the path. As a result, the total amount received at the destination can deviate substantially from a mid-market benchmark due to hidden markups, uncertain fees, and routing choices that prioritize bank convenience over user value.

Recent reports highlight that international students and families experience significant financial losses as a result of these FX gaps and hidden costs. For example, Indian families alone lost an estimated \$200 million in 2024 due to unfavorable exchange-rate markups and banking fees, delaying tuition payments and increasing financial strain [1]. Similar dynamics affect Latin American students and remitters using traditional correspondent banking corridors, where spreads and fees are often opaque and corridor-specific.

Modern remittance services such as Wise improve transparency by offering mid-market rates and explicit fees. However, even these modern platforms conceptually compute the transfer as a single direct conversion, e.g., USD $\rightarrow$ BRL. They do not explicitly optimize over multi-hop institution-aware paths such as USD $\rightarrow$ EUR at one bank, EUR $\rightarrow$ PYG at a second bank, and PYG $\rightarrow$ BRL at a third bank, even if such compositions could yield a higher final amount under realistic spreads and corridor heterogeneity. In practice, FX markets are segmented: different institutions quote different spreads on different corridors, maintain different fee structures, and support different currency sets. When these differences are combined with volatility, a route that looks longer in graph distance can be economically superior.

Formally, given:

- a directed network of banks and currencies,
- a transaction directive specifying a source bank i_s, source currency c_s, destination bank i_t, destination currency c_t, and initial amount a_0, and
- a stochastic model of commercial FX rates and institution-specific spreads and fees,

the central problem is to design a routing policy that selects a sequence of actions (transfers and FX conversions) that begins at (i_s, c_s) with amount a_0, ends at (i_t, c_t) with some random final amount A_T, and maximizes an appropriate functional of A_T (e.g., expected value or a risk-aware criterion) subject to feasibility constraints. Our work proposes and analyzes a particular greedy stochastic policy for this problem and evaluates its performance relative to a commercial benchmark.

III. Methodology [2]

We now formalize the institution–currency network, the stochastic FX environment, and the greedy routing policy implemented in our simulator. All experiments are conducted over a network of twelve institutions spanning four currencies: US Dollars (USD), Brazilian Reais (BRL), Paraguayan Guaraníes (PYG), and Euros (EUR).

A. FX Data and Stochastic Environment

Given access to a historical dataset of daily commercial FX rates among the four currencies:

$$\mathcal{D} = \{R^{\text{comm}}_{x\to y}(t) : x, y \in \{\text{USD}, \text{BRL}, \text{PYG}, \text{EUR}\}, \\ x \neq y,\, t \in \mathcal{T}\},$$

where $\mathcal{T}$ indexes trading days from 2015 to 2025. In the CSV file of the rates, each row corresponds to one date t and contains columns of the form x_y, denoting the commercial rate "units of y per one unit of x" on that day. For any ordered pair (x, y), if the direct column x_y is missing we use the inverse of y_x:

$$R^{\text{comm}}_{x\to y}(t) = \begin{cases} \text{rate}[x\_y](t) & \text{If column exists} \\ 1/\text{rate}[y\_x](t) & \text{Otherwise} \end{cases}$$

We make the following modeling choice for the stochastic environment:

- Each simulation run $n = 1, \ldots, N$ samples one date $\tau^{(n)}$ uniformly from $\mathcal{T}$. Conditional on $\tau^{(n)}$, the commercial rates $R^{\text{comm}}_{x\to y}$ are treated as fixed for the entire routing episode. This represents one plausible "tomorrow" FX snapshot.
- Institution-specific spreads around commercial are modeled as Gaussian in basis points (bps). For each institution i and FX pair (c, c') that it supports internally, we specify a mean markup $\mu_{i,(c,c')}$ and standard deviation $\sigma_{i,(c,c')}$ in bps. For any conversion at time step t of run n, we sample

  $$M^{(n,t)}_{i,(c,c')} \sim \mathcal{N}(\mu_{i,(c,c')}, \sigma^2_{i,(c,c')}),$$

  and define the institution's effective rate as

  $$R^{\text{eff}}_{i,(c\to c')}(\tau^{(n)}, t) = R^{\text{comm}}_{c\to c'}(\tau^{(n)}) \left(1 + \frac{M^{(n,t)}_{i,(c,c')}}{10{,}000}\right).$$

- Each institution i also charges a percentage FX fee $\phi_{i,(c,c')}$ on the gross converted amount and a flat outbound transfer fee $F_{i,c}$ in currency c for same-currency transfers.

The amounts are discretized using a currency-specific rounding operator that reflects real-world monetary precision. Each currency c is assigned a decimal precision d_c: USD, EUR, and BRL are quoted to two decimal places because they have cent-based subunits, while PYG is effectively transacted in whole units and is therefore modeled with zero decimals. Formally,:

$$d_c = \begin{cases} 2 & \text{If } c \in \{\text{USD, BRL, EUR}\} \\ 0 & \text{If } c = \{PYG\} \end{cases}$$

Given a raw amount a in currency c, the model applies:

$$\text{round}_c(a) = \text{round}(a, d_c)$$

after every fee or conversion.

B. Institution: Currency Network

Let $\mathcal{I}$ denote the set of institutions (banks) and $\mathcal{C}$ the set of currencies. Each institution $i \in \mathcal{I}$ is defined by:

- a set of account currencies $\mathcal{C}_i \subseteq \mathcal{C}$ in which it can hold customer balances,
- a set of allowed internal FX directions

  $$\mathcal{F}_i \subseteq \{(c, c') \in \mathcal{C}_i \times \mathcal{C}_i : c \neq c'\},$$

- transfer fees $F_{i,c}$ for each $c \in \mathcal{C}_i$,
- FX markup parameters $\mu_{i,(c,c')}$ and $\sigma_{i,(c,c')}$ for $(c, c') \in \mathcal{F}_i$,
- FX percentage fees $\phi_{i,(c,c')}$ for $(c, c') \in \mathcal{F}_i$.

Examples include:

- SFCU_US: a US institution with USD accounts only and no internal FX capability.
- Chase_US: a US bank with USD accounts and internal USD$\leftrightarrow$BRL, USD$\leftrightarrow$EUR FX pairs, but no PYG accounts.
- BancoNacional_PY: a Paraguayan bank with PYG, USD, BRL accounts and PYG$\leftrightarrow$USD, PYG$\leftrightarrow$BRL FX pairs.

- BancoSantander_ES: a European bank with EUR, BRL, PYG accounts and EUR↔BRL, EUR↔PYG, BRL↔PYG FX pairs.

We construct a directed institution–currency graph $G = (V, E)$, where each node is a feasible bank–currency account:

$$V = \{(i, c) : i \in \mathcal{I},\, c \in \mathcal{C}_i\}.$$

Edges come in two types:

1) **Transfer edges** connect different institutions in the same currency. For each currency c we specify a set of allowed transfers
$$\mathcal{T}_c \subseteq \mathcal{I} \times \mathcal{I},$$
and add an edge $((i, c) \to (j, c))$ to E whenever $(i, j) \in \mathcal{T}_c$ and $c \in \mathcal{C}_i \cap \mathcal{C}_j$. Traversing this edge deducts the flat fee $F_{i,c}$ from the amount.
2) **FX edges** connect different currencies within the same institution. For each $(i, (c, c')) \in \mathcal{F}_i$ we add an edge $((i, c) \to (i, c'))$ to E. Traversing this edge applies the stochastic effective rate $R^{\text{eff}}_{i,(c \to c')}$ and percentage fee $\phi_{i,(c,c')}$ on the gross converted amount.

We use this graph both for routing and for reachability checks: an action that leads to a node from which the destination (i_t, c_t) is not reachable in G is discarded.

C. State, Actions, and Dynamics

We model the routing episode for a given directive as a finite-horizon stochastic process. A user specifies

$$\text{directive } d = (i_s, c_s, i_t, c_t, a_0),$$

where i_s is the source bank, c_s the source currency, i_t the destination bank, c_t the destination currency, and a_0 the initial amount. The initial state is

$$s_0 = (i_0, c_0, a_0) = (i_s, c_s, a_0).$$

At discrete time $t = 0, 1, \ldots$, the state is

$$s_t = (i_t, c_t, a_t) \in \mathcal{I} \times \mathcal{C} \times \mathbb{R}_{\geq 0},$$

and we allow two types of actions:

a) Transfer action.: If $(i_t, j) \in \mathcal{T}_{c_t}$ and both i_t and j support c_t, we may select an action that transfers funds from (i_t, c_t) to (j, c_t). The dynamics are

$$a_{t+1} = \text{round}_{c_t}\left(\max(a_t - F_{i_t,c_t}, 0)\right), \quad (i_{t+1}, c_{t+1}) = (j, c_t).$$

b) FX action.: If $(c_t, c') \in \mathcal{F}_{i_t}$ for some c', we may select an FX action at institution i_t converting $c_t \to c'$. Given the sampled commercial row $\tau^{(n)}$ and markup $M^{(n,t)}_{i_t,(c_t,c')}$, we compute

$$R^{\text{eff}}_{i_t,(c_t \to c')} = R^{\text{comm}}_{c_t \to c'}(\tau^{(n)}) \left(1 + \frac{M^{(n,t)}_{i_t,(c_t,c')}}{10{,}000}\right),$$

$$\tilde{a}_{t+1} = a_t \cdot R^{\text{eff}}_{i_t,(c_t \to c')}, \quad \text{fee}^{\text{FX}}_t = \phi_{i_t,(c_t,c')} \cdot \tilde{a}_{t+1},$$

$$a_{t+1} = \text{round}_{c'}\left(\max(\tilde{a}_{t+1} - \text{fee}^{\text{FX}}_t, 0)\right), \quad (i_{t+1}, c_{t+1}) = (i_t, c').$$

The episode terminates when one of the following occurs:

- the destination node is reached: $(i_t, c_t) = (i_t, c_t)$,
- a maximum hop limit $H_{\max}$ is exceeded,
- the available amount a_t becomes non-positive,
- no feasible actions remain (dead end).

We also track visited nodes (i, c) and disallow returning to them to avoid trivial cycles.

D. Greedy One-Step Lookahead Policy

We adopt a myopic greedy policy with one-step Monte Carlo lookahead.

At state $s_t = (i_t, c_t, a_t)$ in run n:

1) We enumerate all feasible transfer and FX actions from s_t that
   - respect the institution's capabilities (account currencies and FX pairs),
   - do not return to a visited node,
   - lead to a node (i', c') from which (i_t, c_t) is reachable in G.
2) For each candidate action u, we estimate the expected next-state amount $\mathbb{E}[a_{t+1} \mid s_t, u]$ using K-Monte Carlo samples. Transfers are deterministic, so this expectation is just the post-fee amount. FX actions involve sampling $M^{(n,t,k)}_{i,(c,c')}$ for $k = 1, \ldots, K$ and averaging the resulting $a^{(k)}_{t+1}$.
3) To compare actions across different next currencies, we normalize expected amounts into destination currency c_t using the commercial rate from the sampled row:
$$\tilde{a}^{\text{dest}}(u) = \mathbb{E}[a_{t+1} \mid s_t, u] \cdot R^{\text{comm}}_{c' \to c_t}(\tau^{(n)}),$$
where c' is the next currency under u.
4) We add a simple graph-distance bonus to encourage progress toward the destination. Let d_t be the shortest-path distance (in hops) from the current node (i_t, c_t) to (i_t, c_t) in G, and $d_{t+1}(u)$ the distance from the candidate next node. We define a scalar score
$$\text{score}(u) = \tilde{a}^{\text{dest}}(u) \cdot (1 + \alpha\,(d_t - d_{t+1}(u))),$$
where $\alpha > 0$ controls how strongly we reward getting closer in graph distance. In our implementation we set $\alpha = 0.10$, so reducing the distance by one hop yields approximately a 10% bonus in the score.
5) The greedy policy selects
$$u^\star_t = \arg\max_u \text{score}(u),$$
and then executes a single stochastic realization of that action (drawing a fresh markup for FX if applicable) to update the state from s_t to s_{t+1}.

This yields a simple, interpretable policy that balances two objectives: maximizing expected immediate value in the destination currency and making structural progress towards the receiver in the institution–currency graph.

E. Benchmark and Fairness Ratio

For each simulation run n and directive d, we define a commercial benchmark that represents an idealized direct conversion at mid-market rates with no institution-specific markups or fees:

$$B^{(n)} = a_0 \cdot R^{\text{comm}}_{c_s \to c_t}(\tau^{(n)}).$$

This is the amount the recipient would receive if the entire initial amount were converted directly from c_s to c_t at the sampled commercial rate for that day and no additional costs were incurred.

Let $A_T^{(n)}$ be the final amount at the destination node (i_t, c_t) under the greedy policy in run n, and define $A_T^{(n)} = 0$ if the route fails to reach the destination within the hop or feasibility constraints. For successful runs, we define the *fairness ratio*

$$\rho^{(n)} = \frac{A_T^{(n)}}{B^{(n)}}.$$

A value $\rho^{(n)} > 1$ indicates that the realized multi-hop route delivered more than the commercial benchmark, while $\rho^{(n)} < 1$ reflects a loss relative to that idealized direct conversion.

Across N runs, we summarize performance using:

- the success rate $\hat{p}_{\text{succ}} = \frac{1}{N} \sum_n \mathbf{1}\{A_T^{(n)} > 0\}$,
- the mean and standard deviation of $A_T^{(n)}$ conditional on success,
- the mean and standard deviation of $\rho^{(n)}$ conditional on success,
- and the fraction of successful runs with $\rho^{(n)} \geq 1$.

These statistics provide a distributional view of how the greedy policy performs under different FX scenarios drawn from historical data.

F. Model Assumptions and Limitations

Our modeling choices involve several simplifying assumptions:

- i.i.d. FX snapshots: We treat each simulation run as an independent draw of the entire FX surface from historical data. In reality, FX paths over time are correlated and transfers may span multiple days; modeling temporal dynamics would require a more complex Markovian or stochastic-process model.
- Gaussian markups: Institution spreads around commercial are modeled as Gaussian in basis points. This is convenient but may underrepresent tail behavior or asymmetric pricing strategies.
- Static network: The institution–currency graph and transfer relationships are fixed. We do not model dynamic link failures, capacity constraints, or time-varying corridor availability.
- Myopic policy: The greedy policy optimizes only one step ahead. It may miss routes where a short-term sacrifice (e.g., going through a high-fee edge) enables a much more favorable conversion later.
- No Time or Risk Preferences: We ignore transfer times, settlement risk, and user-specific risk aversion. The objective is purely expressed in terms of final amount in the destination currency and a relative benchmark.

Despite these limitations, the model captures important real-world features that simpler currency-only or deterministic frameworks miss: heterogeneity across institutions, corridor-specific spreads, explicit fee structures, and the combinatorial structure of routing over bank networks.

IV. Analysis and Results

Fig. 1 shows the full institution–currency network produced by the model, represented as an undirected, acyclic graph containing 31 nodes and approximately 60 total edges. This graph encodes all feasible transfer and FX relationships and forms the structural basis for routing decisions. Using this network,

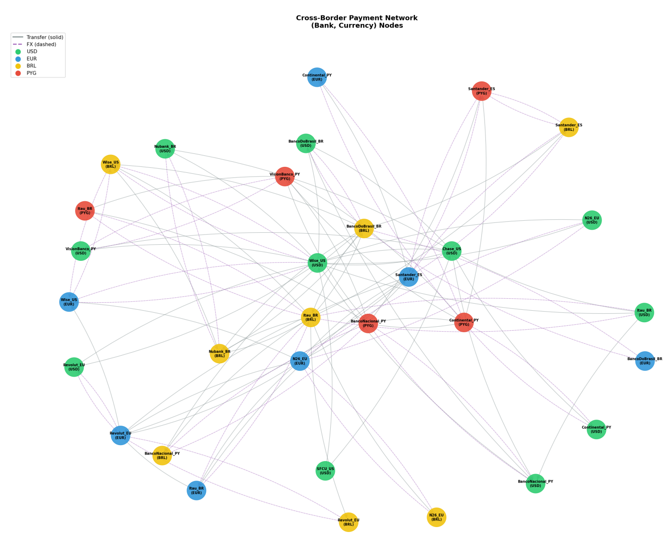

Fig. 1. Graph of all nodes and edges with states, actions

the model successfully solves the directive of sending $10,000$ BRL from Santander_ES to Revolut_EU in EUR. Across 500 Monte Carlo simulations using ten years of historical FX data, the routing algorithm achieves a 100% success rate, as stated in Listing 1, which shows the path cost breakdown. The highlighted path in Fig. 2 shows the optimal path, where the purple node is the starting node (Santander_ES in BRL) and the orange one is the destination node(Rveoult_EU in EUR).

Listing 1. Santander_ES BRL $\rightarrow$ Revoult_EU EUR Path Cost Breakdown

```
Directive: 10,000.00 BRL from Santander_ES
        to: EUR at Revolut_EU
Commercial Benchmark: 1,664.54 EUR
Actual Received:      1,867.10 EUR
Fairness Ratio:       1.1217

START: (Santander_ES, BRL)
  Amount: 10,000.00
  Step 1:
```

```
    FX at Santander_ES: BRL -> PYG
    Amount: 10,000.00 -> 14,559,723.00
    (rate: 1462.994642, fee: 70,223.74)
  Step 2:
    FX at Santander_ES: PYG -> EUR
    Amount: 14,559,723.00 -> 1,871.10
    (rate: 0.000129, fee: 9.40)
  Step 3:
    TRANSFER Santander_ES ->
    Revolut_EU [EUR]
    Amount: 1,871.10 -> 1,867.10 (fee: 4.00)
END: (Revolut_EU, EUR) Amount: 1,867.10
```

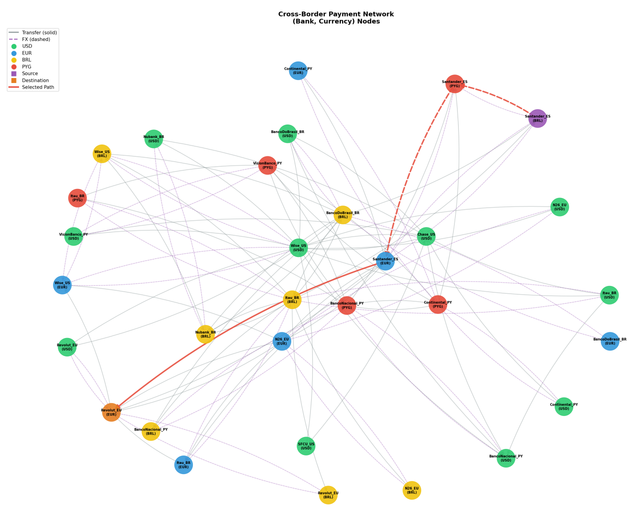

Fig. 2. Santander_ES BRL → Revoult_EU EUR Best Path

The resulting distributions in Fig. 3 illustrate how FX volatility shapes outcomes: the final-amount distribution (left) displays a bimodal pattern driven by variation in historical BRL/EUR rates, while the fairness-ratio distribution (right) is tightly concentrated around a mean of 1.13, indicating consistent outperformance of the commercial benchmark. The model delivers a mean of $2,022.83$ EUR against a benchmark of $1,664.54$ EUR, demonstrating substantial added value from institution-aware multi-hop routing.

The model also correctly identifies scenarios in which no valid routing path exists. For example, the directive Nubank_BR → SFCU_US fails systematically because SFCU_US functions as a send-only node with no inbound USD edges, while Nubank_BR has no outgoing USD transfer relationships, making SFCU_US unreachable regardless of FX conditions. This structural asymmetry mirrors realistic banking constraints, where certain institutions can initiate but not receive international transfers, or vice versa. Finally, although many corridors yield fairness ratios above one, the model also captures cases in which the routed outcome underperforms the benchmark. The distributional behavior and path diagnostics for those scenarios are provided in the Appendix.

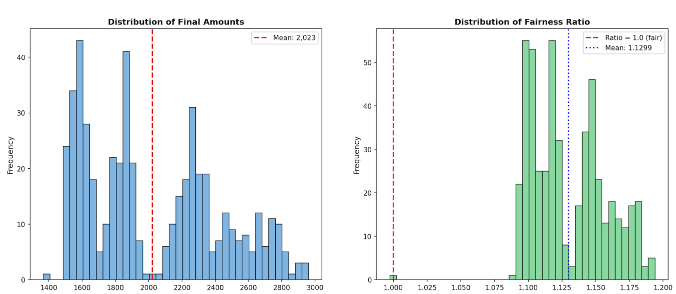

Fig. 3. Santander_ES BRL → Revoult_EU EUR Distribution

V. Alternative Approaches

Beyond the greedy step-wise optimization model implemented in this work, we examined two additional approaches: a naive fixed-path evaluation and a full POMDP formulation. Both are conceptually useful baselines, but each is ultimately less suitable than our chosen model for the goals and constraints of this project.

A. Naive Fixed-Path Evaluation

The naive model treats the cross-border payment network as if all exchange rates, spreads, and fees were deterministic and time-invariant. Each edge in the institution–currency graph is assigned a fixed effective rate and fee based on historical averages or representative corridor values. Once G is endowed with these static weights, the optimal route can be approximated by applying a standard shortest-path algorithm, like Dijkstra, on a suitably defined cost function (such as negative log effective rate or total fee).

This approach is computationally simple and helps build intuition for how network structure affects routing. However, it removes the key source of difficulty in real-world routing: uncertainty. By freezing FX rates and conversion costs, the naive model cannot capture situations where a multi-hop route becomes temporarily advantageous due to stochastic fluctuations in spreads or where risk considerations (e.g., avoiding exposure to volatile corridors) should shape the chosen path. It also fails to represent the variability that international users actually experience. The naive model serves as a conceptual baseline, but the institution-aware greedy policy is better aligned with the stochastic nature of FX markets.

B. Full POMDP Formulation

A fully realistic formulation of cross-border routing is naturally a POMDP, because the agent does not observe the true future evolution of exchange rates during the transfer. In practice, the sender only sees indicative quotes or delayed information, while the actual conversion may occur later at a different rate. This delayed settlement means the agent acts under uncertainty about the FX environment that will prevail when each conversion occurs.

The latent state underlying the routing problem includes:

- the commercial FX surface at the (unknown) execution time of each hop,
- institution-specific markups drawn at conversion time,

- any corridor-specific frictions not observable ex ante.

Because these components evolve stochastically and are not observed at decision time, the agent should in principle maintain a belief distribution over future FX realizations and select actions that maximize expected value under that belief.

However, tackling the full POMDP directly would require jointly modeling belief updates, real-time FX dynamics, and institution-specific behaviors, making it beyond complex given the time and data constraints for the project. Instead, we adopt a divide-and-conquer approach by isolating and studying a tractable core subproblem—routing under uncertainty generated by stochastic FX snapshots and institution-specific spreads. By first understanding this simplified but realistic component, we gain insight into the broader routing question without committing to the full complexity of a general POMDP solution.

VI. Further Research

The greedy policy and institution–currency model presented here are a first step toward realistic stochastic routing for cross-border payments. Several extensions are natural directions for future work.

First, the network could be expanded to include more currencies and institutions, especially those relevant to major remittance corridors in Africa, Asia, and Latin America. This would allow us to study how network density and corridor coverage influence the prevalence of advantageous multi-hop routes.

Second, the stochastic environment could be enriched by modeling FX as a time-series process rather than i.i.d. snapshots, incorporating correlations across currency pairs, and using more realistic (possibly heavy-tailed or skewed) distributions for institution spreads. This would enable analysis of time-sensitive strategies that condition routing decisions on current volatility regimes.

Third, alternative routing policies could be explored. Examples include limited-horizon lookahead (e.g., two- or three-step Monte Carlo tree search), risk-sensitive objectives that penalize downside variance in A_T, and learning-based policies that adapt to observed outcomes over time. These methods would move the framework closer to reinforcement learning or approximate dynamic programming over financial networks.

Fourth, the fairness metric itself could be generalized. In addition to the simple ratio of final amount to commercial benchmark, one could consider user-centric criteria such as minimum guaranteed amount with a specified confidence level, or utilities that trade off expected value against variability and delay.

Finally, the project could be extended toward a more complete POMDP formulation by explicitly modeling information structures (e.g., quote updates, corridor outages) and by learning institutional parameters from real transaction data rather than hand-crafted priors. Such extensions would push the model closer to deployment in decision-support tools for remittance platforms or banks, while maintaining the core insight of this work: routing decisions in cross-border payments are fundamentally stochastic and network-aware, and simple greedy policies over an institution–currency graph already reveal meaningful opportunities for improving the value delivered to end users.

VII. Conclusion

In this project we introduced a greedy stochastic routing framework for cross-border payments over an institution–currency network and showed that even a simple one-step lookahead policy can exploit heterogeneity in spreads, fees, and corridor structure. Using ten years of historical FX data and institution-specific markups, the model constructs feasible multi-hop routes, evaluates them via Monte Carlo simulations, and benchmarks outcomes against a commercial mid-market conversion. For the realistic directives we study, the policy attains a 100% routing success rate, consistently finds short three-hop paths, and delivers final amounts that systematically exceed the benchmark. At the same time, the model correctly identifies structural failure cases where network asymmetries make the destination unreachable, as well as cases where the final amount is maximized but still lies below the benchmark (see Appendix).

While the framework abstracts away full FX dynamics, settlement risk, and user-specific preferences, it already captures key real-world features that standard currency-only or deterministic models miss. The results suggest that institution-aware routing over bank networks can meaningfully improve value delivered to recipients, even under noisy and volatile FX conditions. Extending this approach to richer stochastic environments, larger networks, and more sophisticated routing policies is a promising direction for future work and could inform practical decision-support tools for remittance platforms and banks.

Appendix

Fig. 4 shows the highlighted three-step path for the directive SFCU_US (USD) to Continental_PY (PYG), where SFCU_US initiates a transfer of $1,000$ USD, Chase_US acts as the sole feasible USD correspondent, and Continental_PY performs the final USD→PYG conversion. Listing 2 confirms this structure numerically: the model consistently applies two USD transfers followed by a single FX step, yielding a final amount of roughly 7.3 M PYG against a 7.4 M PYG benchmark. The distributions in Fig. 5 mirror this behavior. The final-amount distribution (left) is bimodal due to historical USD/BRL/EUR volatility, whereas the fairness-ratio distribution (right) is sharply concentrated around 0.986, indicating stable relative performance despite FX noise. This path exemplifies a case where routing succeeds reliably but the final value remains slightly below the commercial benchmark (fairness ratio < 1).

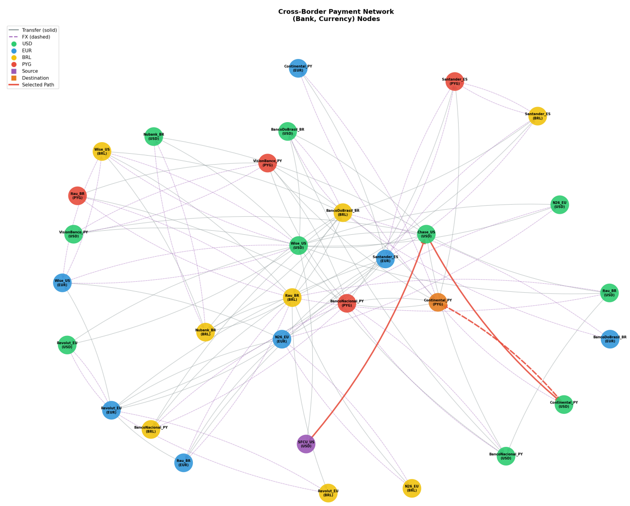

Fig. 4. SFCU_US USD → Continental_PY PYG Best Path

Listing 2. SFCU_US USD → Continental_PY PYG Path Cost Breakdown

```
Directive: 1,000.00 USD from SFCU_US
         to: PYG at Continental_PY
Commercial Benchmark: 7,401,971.50 PYG
Actual Received:      7,298,222.00 PYG
Fairness Ratio:       0.9860
START: (SFCU_US, USD) Amount: 1,000.00
  Step 1:
    TRANSFER SFCU_US -> Chase_US [USD]
    Amount: 1,000.00 -> 995.00 (fee: 5.00)
  Step 2:
    TRANSFER Chase_US ->
    Continental_PY [USD]
    Amount: 995.00 -> 985.00 (fee: 10.00)
  Step 3:
    FX at Continental_PY: USD -> PYG
    Amount: 985.00 -> 7,298,222.00
    (rate: 7427.932550, fee: 18,291.28)
END: (Continental_PY, PYG)
Amount: 7,298,222.00
```

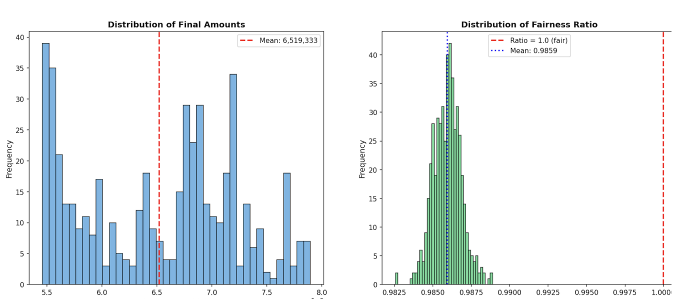

Fig. 5. SFCU_US → Continental_PY PYG Distribution

Another simulation test is Fig. 6, where we transfer 10,000,000 PYG from Banco Nacional to Chase in USD.

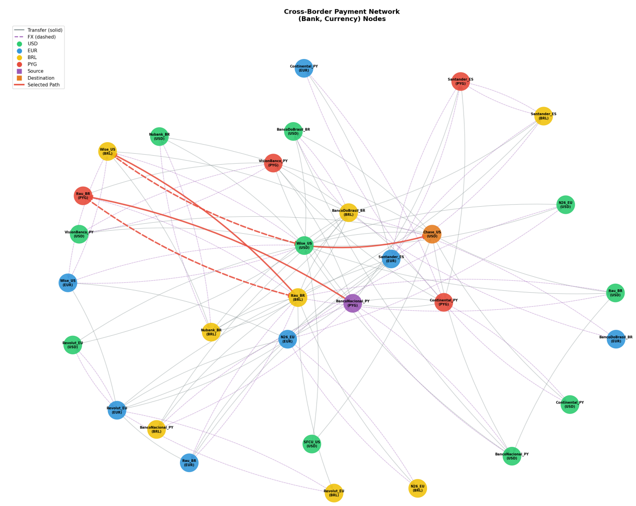

Fig. 6. BancoNacional_PY PYG → Chase_US USD Best Path

According to Listing 3, across all stochastic rate realizations, the mean final amount is only 1,299 USD, with a fairness ratio averaging 0.8558, meaning the system delivers around 14–15% less than the commercial benchmark on average.

Listing 3. BancoNacional_PY PYG → Chasel_US USD Path Cost Breakdown

```
Directive:
10,000,000.00 PYG from BancoNacional_PY
to: USD at Chase_US
Commercial Benchmark: 1,380.70 USD
Actual Received: 1,197.81 USD
Fairness Ratio: 0.8675
START: (BancoNacional_PY, PYG)
Amount: 10,000,000.00
Step 1:
    TRANSFER BancoNacional_PY -> Itau_BR [PYG]
Amount:
    10,000,000.00 -> 9,985,000.00
    (fee: 15,000.00)
Step 2:
    FX at Itau_BR: PYG -> BRL
    Amount: 9,985,000.00 -> 6,769.41
    (rate: 0.000681, fee: 27.19)
Step 3:
    TRANSFER Itau_BR -> Wise_US [BRL]
    Amount: 6,769.41 -> 6,764.41 (fee: 5.00)
Step 4:
    FX at Wise_US: BRL -> USD
    Amount: 6,764.41 -> 1,199.81
    (rate: 0.178083, fee: 4.82)
Step 5: TRANSFER Wise_US -> Chase_US [USD]
    Amount: 1,199.81 -> 1,197.81 (fee: 2.00)
END: (Chase_US, USD) Amount: 1,197.81
```

Fig. 7 reports that the final-amount distribution (left) has a mean of about $1,299$ USD, with outcomes spread roughly between $1,100$ USD and $1,500$ USD, reflecting volatility in the underlying PYG/USD corridor. The fairness-ratio distribution (right) is tightly concentrated around a mean of 0.8558, well below one, highlighting that the multi-hop routes available to Banco Nacional simply cannot match the direct-market equivalent.

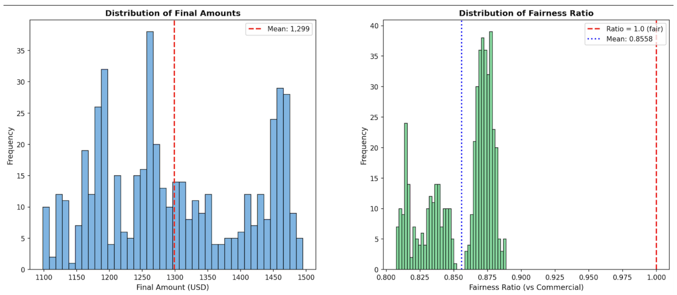

Fig. 7. BancoNacional_PY PYG → Chasel_US USD Distribution

ACKNOWLEDGMENT

We used ChatGPT to develop the initial coding framework for the routing model, including the data structures, graph construction, and simulation pipeline. Building on that foundation, we extrapolated the FX data, implemented the stochastic environment, and performed all necessary computations. Claude was later used to refine and finalize the full codebase, ensuring consistency across modules and resolving implementation details.

REFERENCES

[1] S. Sankar, "International Students Face Major Financial Losses Due to Forex Gaps", VisaVerge, August. 2025. https://www.visaverge.com/knowledge/international-students-face-major-financial-losses-due-to-forex-gaps/ (accessed Dec. 08, 2025).

[2] M. K. Kochenderfer, T. A. Wheeler, and K. H. Wray, " Algorithms for Decision Making", MIT, 2022

TEAM MEMBERS CONTRIBUTION

Marco Paes (4 units: Brainstorming; Research, Model Computation & Data Analysis; Data Extraction, Coding; Edit Paper. Interviewed international students to understand the practical challenges they face when sending money across borders, as well as with financial-industry and fintech professionals to validate our understanding of current systems, difficulties, and emerging technologies. These conversations helped (1) refine the scope and direction of the model so that it could meaningfully address real user pain points, and (2) select data sources and guide the computations, analysis, and simulations, ensuring that we included the elements most relevant to the problem and the purpose of building a base model.

Andrea Nam (3 units): Brainstorming, Model Computation & Data Analysis; Organize Ideas; Write Paper; Debugging.