An Assessment Framework to Offline Functionality in Central Bank Digital Currencies

Vladyslav Nekriach*, Panagiotis Michalopoulos*, Cyrus Minwalla†, Andreas Veneris*‡
* Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada
‡ Department of Computer Science, University of Toronto
{p.michalopoulos, vladyslav.nekriach}@mail.utoronto.ca, veneris@eecg.toronto.edu
† Bank of Canada, Ottawa, Canada
CMinwalla@bank-banque-canada.ca

***Abstract*—The emergence of blockchain technology has reshaped the financial sector, leading to the introduction of new financial products and services. In parallel, central banks worldwide are investigating the possibility of issuing central bank digital currency (CBDC) and the potential use of blockchain as a foundational technology. Of particular interest is offline functionality, where transactions can be completed even if neither party is connected to the ledger at the time of transaction. Such offline functionality carries unique opportunities for countries and policy makers, but also comes with a new set of risks. Proposed herein is an assessment framework as a critical modeling tool towards risk evaluation. The framework is broadly applicable to blockchain-based solutions that support multiple off-chain transactions prior to synchronization. Urban and rural environments were modeled to demonstrate the model's fidelity and flexibility. A test offline digital currency solution was evaluated via threat level experiments to ascertain the prototype's resilience to varying levels of malicious activity. Results illustrate how various parameter configurations affect resilience to malicious activity and information propagation, demonstrating the effectiveness of the framework in providing valuable insights for the design of offline currency systems.**

***Index Terms*—digital currency, blockchain, offline, CBDC, simulation, agent-based, evaluation framework**

I. Introduction

The advent of blockchain technology has sparked the emergence of a vast ecosystem of digital currencies with over $113B in total assets invested in this space [1]. The technology has fundamentally transformed our understanding of, and relationship to, money. This revolution has transcended its origins within Decentralized Finance and speculative trading communities, initiating significant innovation within traditional financial institutions as evidenced by the introduction of cryptocurrency Exchange-Traded Funds (ETFs) [2], institutional custody solutions, and most notably, the development of *Central Bank Digital Currencies (CBDCs)*. These national currencies are driven by the requirements of end-users, namely consumers and merchants. As part of their investigation, central banks are exploring the potential uses of blockchain and distributed ledger technologies as the settlement layer for their own versions of digital currencies [3]–[5]. Such blockchain-backed CBDCs would share many common characteristics with other cryptocurrencies, including those existing in decentralized finance and stablecoins.

Past experiments in this space confirm that CBDCs have the potential to provide benefits in jurisdictions where access to digital money or the digital infrastructure may be inadequate. One pivotal aspect is their ability to function in the absence of network connectivity, referred to as *offline* mode. In this offline mode, users can transact with each other without requiring a connection to the online ledger of the bank [6]. Such a mode would increase access in remote populations, areas with unreliable connectivity, and during infrastructure failures [6], [7]. If offered as a stand-alone product without requiring a bank account, it could improve access for the unbanked, tourists and children. While offline use carries multiple benefits, it also introduces novel challenges and comes with security risks. Understanding which risks are relevant among the gamut of potential options, and how scope and scale can change the impact, are key steps in making informed policy decisions around CBDC.

To that end, this paper presents an assessment framework as a modeling tool for assessing risk. An institution may use it to evaluate an offline CBDC system under different operating conditions and when exposed to key failure modes such as double-spending attacks and fraudulent transactions. Such a framework informs important questions on offline digital currency designs, such as the length of time funds can stay disconnected from a central ledger, impact of various attacks on overall system health, and the effectiveness of strategies to detect and neutralize bad actors.

The framework can be adjusted to emulate a wide range of use cases from urban settings with dense transaction activity to rural settings characterized by sparse connectivity. Nodes in this network represent geographic economic zones containing one or more agents (*e.g.,* users, merchants) that transact with each other, with agents being capable of moving between zones. The Briolette project [8], a token-based payment system, was used as the test system due to its open source nature and built-in support for offline transactions. The suitability of the proposed framework is demonstrated through experimental evaluations and the resulting insights.

The remainder of this paper is organized as follows: Section II introduces the necessary background regarding CBDCs and related work. Section III presents the proposed assessment framework, its parameters and overall architecture. Section IV presents Briolette and its evaluation using the framework, and Section V concludes the paper.

II. Background and Related Work

A central bank digital currency (CBDC) is a digital form of money issued by the central bank [9]. Once issued, CBDCs become a direct liability of the central bank, distinguished

from physical banknotes and deposits of commercial banks held by the central bank. The difference lies in being recorded separately in purely digital form [10], [11]. CBDCs are generally categorized into two types: *wholesale* and *retail*. Wholesale CBDCs are designed for large-value transactions between financial institutions, while retail CBDCs are made available directly to individuals and businesses for everyday payments, functioning as a digital analogue of cash. In this work, we focus specifically on retail CBDCs with support for offline functionality.

CBDCs can also be divided into *token-based* and *account-based* structures [12], [13] based on how users access the currency. Token-based systems rely on the exchange of cryptographic tokens with a pre-assigned value, while account-based ones rely on balances and some form of identity verification. However, certain proposals [14] have emerged that challenge this dichotomy by incorporating elements of both paradigms, thus blurring the distinction between token-based and account-based CBDC models.

Offline CBDCs are a special case as they operate without network connectivity or access to an online ledger [6]. Several design choice arise with regard to transaction settlement and the duration that a wallet can remain offline. The funds representation could be account-based or token-based, each with its own set of risks. If settlement occurs offline, then users can re-spend received funds in subsequent offline transactions without needing online synchronization. Similarly, the length of time a CBDC wallet can remain offline must be balanced with the need for periodic synchronizations to the ledger. The Bank of International Settlements distinguishes between the offline CBDC types [6] with the most practical being *intermittently offline*, where funds can be re-spent consecutively but a synchronization period is required.

The high stakes associated with the CBDC deployment have contributed to a cautious stance among central banks with regard to issuing a CBDC [15]. Ergo, despite significant interest and proposed designs for *offline CBDCs* [16], [17], real-world implementations remain scarce [18]. This lack of practical implementations poses a challenge, as it limits the ability to test designs in the field, identify potential weaknesses, and iteratively improve upon them. Simulations are an attractive alternative to an in-field evaluation. They can yield insights on security, privacy and performance under various operating conditions, consumer behaviours and socio-economic factors. Furthermore, certain failure modes can be emulated to understand/assess their impact in a controlled environment without risking substantial financial harm [19]–[22].

Multiple works have simulated CBDC designs, primarily focusing on how they augment the existing financial system [20], [23] and on user adoption [21]. However, these studies either briefly touch upon the offline functionality or omit it entirely. The authors in [22] simulated an offline CBDC system in the context of its native region of Norway. The authors use the Barabási-Albert model to generate the topology of the network, but in contrast to this paper, nodes are used to represent users of the system and the spatiotemporal movement of the agents is simulated by setting a high average number of connections per node. The protocol was evaluated against malicious users and tested multiple mitigation mechanisms.

TABLE I: Framework Hyperparameters

Category	**Hyperparameters**
Deployment scenario	Graph structure Number of merchants Number of banks
Information propagation	Peer to merchant probability Peer to peer probability Peer to bank probability Peer movement probability
Risk control and security	Withdrawal amount Merchant synchronization frequency Number of offline transactions allowed Ratio of honest to malicious agents

The proposed framework is developed with a focus on token-based offline CBDC systems, as account-based designs introduce a different set of risks that require separate consideration. While both models must contend with the possibility of double-spending in the absence of immediate synchronization with the central ledger, the mechanisms and vulnerabilities differ. For example, in token-based systems, control derives from possession of cryptographic keys, which exposes users to irreversible loss in the event of key compromise and increases the difficulty of preventing double-spending during offline transfers. In account-based systems, by contrast, ownership is tied to centralized account records; when those records are unavailable offline, the safeguards against double-spending that normally follow from real-time ledger validation are suspended, leaving the model reliant on hardware enforcement or transaction limits. These divergences highlight that an assessment framework tailored to token-based designs cannot be straightforwardly applied to account-based offline architectures.

III. An Assessment Framework

A framework to assess the performance characteristics and security attributes of offline CBDC systems is presented. The framework is built on top of hybrid network models which can be adjusted through *hyperparameters* to a variety of use-cases. Those use-cases range from a dense urban environment to a rural one, where the communications infrastructure may be sporadic, minimal or absent altogether. This allows for simulation parameters, such as geographic density and agent behaviours, to be tailored to specific jurisdictions. The generated networks can scale to the available compute capability. Table I gives an overview of those hyperparameters as set out below:

1. *Deployment scenario:* The specifics of the region where the CBDC will be deployed, including, but not limited to, how well-connected the locations are with each other, the number of agents, *etc.*
2. *Information propagation:* The description of how information flows through the network. This subset of hyperparameters is typically described in terms of interaction probabilities between network agents and probability of an agent moving through the network.
3. *Risk control and security:* Parameters related to the measures taken by the CBDC system to minimize risk exposure

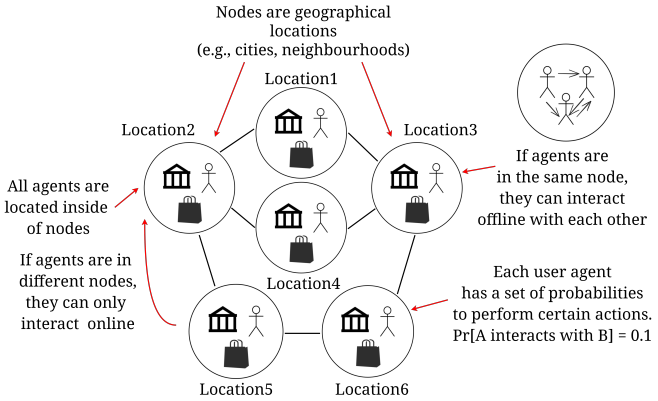

Fig. 1: Overview of the simulation graph structure.

to the adversarial model. Examples include limits on the withdrawal amounts and/or the number of offline transactions before online synchronization is required, ratio of malicious to honest agents, *etc.*

A. Network model

Figure 1 presents the network model $M = (G, A)$ that the framework uses. It is based on a graph $G = (V, E)$ used as an underlying representation of a geographic area. Each node $v \in V$ represents a geographical location on the map (*e.g.,* neighborhood, city), and each edge $e \in E$ represents a physical path that an agent can take to move between those locations. As such, the higher the degree of a node, the more well-connected is the location, leading to more agents moving there and increasing the transaction activity. In this notation, the set A represents a set of agents, where each agent $a \in A$ resides in one of the graph nodes. Agents can represent different entities in the real world, such as banks, merchants, individuals, etc. Agents located in the same node are considered to be within proximity to conduct offline transactions. If an agent needs to interact with another agent outside its node, they would either move to that node using edges from E, or go online. The motivation for such a network model is that it depicts real-life population behavior that lends itself well for use in a simulation.

To model urban scenarios, we employ graphs generated using the Barabási–Albert (BA) model [24]. Urban environments are typically characterized by the presence of *high connectivity hubs*—nodes with significantly more edges—connected to spokes—nodes with lower connectivity. The hubs mimic central areas of a city with increased financial activity, such as the downtown core, while spokes represent areas of lower activity, such as the suburbs. The BA model is well-suited to such a network topology due to its underlying mechanism of preferential attachment, resulting in a *power-law degree distribution*, where the probability $P(m)$ of a node having m connections is proportional to $m^{-\gamma}$, resulting in an uneven distribution of edges, leading to network hubs that mirror those observed in real-world urban systems.

Rural scenarios were modeled with the Watts–Strogatz (WS) network model [25], which generates *small-world networks*—networks that combine high clustering with short average path lengths. Similarly, rural environments are typically marked by localized clustering and limited connectivity beyond these clusters. Therefore, to replicate this structure, the WS model randomly introduces, with a probability p, "shortcut" edges between nodes, thereby significantly reducing the average path length between otherwise distant nodes while preserving high local clustering. The resulting small-world networks exhibit a relatively homogeneous topology, with nodes maintaining approximately uniform degrees, reflecting the more evenly distributed and locally constrained connectivity patterns typical of rural settings.

B. Hyperparameters of interest

1) Deployment scenario: This category of hyperparameters influences the structure of the generated network that represents the deployment topology of the currency system. Our framework considers the *average node degree*, the *presence of hubs*, and the *number of agents* in the network. The first two relate to the underlying graph creation model and influence the information flow in the network, while the last dictates the number of users, merchants, and banks in the simulation.

2) Information propagation: Another important characteristic is how the information flows through the network. If information flow is constrained, system state updates (*e.g.,* revocations) could be delayed indefinitely in systems where updates are propagated through peers. Therefore, we define the following hyperparameters that impact information flow:

Peer to peer/merchant interaction probability: The probability that an agent will participate in a transaction. Controls how the information is propagated between agents inside each node, effectively controlling how fast agents within the same node are notified about the latest state update.

Peer to node movement probability: The probability that an agent will move to a neighbouring node. Controls the information propagation speed through the network, since agents can interact (and propagate system updates) with each other offline only when located in the same node.

Peer to bank probability: Represents the likelihood that an agent accesses their online bank account triggering a direct state update from the bank.

3) Risk control and security: We propose the following hyperparameters that influence the security of a CBDC system to study the potential harm incurred due to malicious attacks:

Maximum withdrawal amount: Influences the time an agent can remain offline by indirectly limiting the number of transactions through the withdrawal amount. Useful when the average transaction amount is large, allowing for more granular control compared to limiting the number of transactions.

Merchant synchronization frequency: Controls how frequently merchants reconcile their offline state with the bank allowing the latter to detect any counterfeit.

Number of offline transactions allowed: Controls how long can an agent be offline in terms of the number of transactions.

Ratio of honest to malicious agents: A percentage of the agents are considered to be malicious users that try to double-spend. Performing sensitivity analysis on this parameter allows to test the level of security that the system can provide.

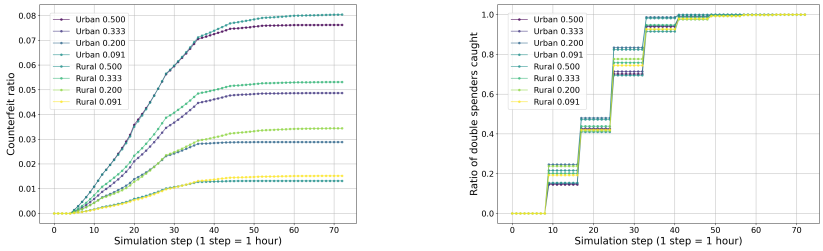

(a) (b)

Fig. 2: Ratio of counterfeit and double spenders caught under different threat scenarios for urban and rural areas.

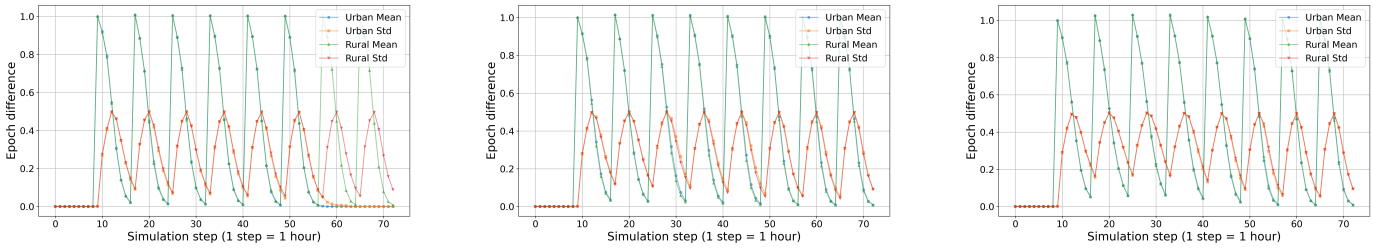

(a) (b) (c)

Fig. 3: Second order statistics of the difference between global and local epochs for threat levels of (a) $\frac{1}{5}$ (b) $\frac{1}{3}$ (c) $\frac{1}{2}$

IV. Experimental Evaluation

A. The Briolette ecosystem

Evaluation of the framework was conducted with Briolette [8] as the test system. Briolette is an open source offline digital currency framework with the aim to facilitate research related to offline currencies. It is a proof of concept system for token-based retail digital currency that supports offline settlement, re-spending of received funds, and token traceability. It was chosen since it embeds a fully functional offline transaction protocol and shares many similarities with a blockchain ecosystem, including the use of digital signatures for proof of ownership, synchronization with a source of truth (consensus) and an epoch denoting block height.

Interaction with the system occurs through a *wallet application*. This wallet must be registered and receive a special credential used to access the various services of the system. The generated credential can be used to obtain tickets from the ticketing service that function as destination addresses and are necessary for sending or receiving tokens. Both tickets and tokens incorporate expiration dates and optionally other transaction limits. Wallets periodically synchronize to a shared state (consensus service) containing the epoch number, certificates for the system's services, and revocation information. Updates to the shared state are provided by the state service and are propagated among participants through a gossip protocol. A lifecycle of a wallet is determined by a delta between the issuance epoch and the current epoch, which is equivalent to the difference in block heights. If a wallet is caught performing malicious behaviour, it is blacklisted and its credentials are revoked the next time it connects for a synchronization.

1) Token lifecycle: A token moves through three phases: mint, transfers, and expiration/revocation. Tokens are assigned a monetary value and an initial recipient—typically, an intermediary Financial Institution, such as commercial banks—before being signed by the mint. During a transaction, tokens are bound to a valid ticket linked to the recipient's credentials via a digital signature, which serves as proof of ownership. Additional transactions generate nested signatures for non-repudiation. Due to this nesting feature, token size grows as a function of transactions, eventually requiring a trimming (re-minting) operation to refresh the token state.

2) Offline transaction: Token transfers between wallets consist of the following steps. First, the two peers validate each other through a gossip handshake that ensures that both share the most recent shared state known to them. Next, the sender will validate the receiver's ticket by validating the signature of an authorized entity and confirming expiry status and any supported attributes. Once successful, the receiver validates the tokens and can choose to accept or reject them. If accepted, the sender binds them to the receiver's ticket and sends them over. It is noted that a correctly constructed token does not protect against duplicates. Thus, singularity of the token can only be established when the token is synchronized with the consensus mechanism. During synchronization, the validation service compares the existing token history with the incoming one. A double-spend event is identified if a fork is detected or the new history is shorter than the existing one. The nested signatures are used to identify the source of the fork and the offending wallet is subsequently blacklisted.

3) State update: The system periodically issues state updates. These updates contain new revocation data obtained by the user wallet revocation service, as well as any new certificates for the services of the system. The updates are signed in order for wallets to be able to check their authenticity.

The proposed framework is not restricted to Briolette. Alternate digital currency systems based on Distributed Ledger Technologies could be considered; for example, already estab-

TABLE II: Experimental Parameters

Parameter	Value
Double spender ratio ($Ratio_{ds}$)	$\frac{1}{11}$, $\frac{1}{5}$, $\frac{1}{3}$, $\frac{1}{2}$ / [$\frac{1}{11}$, $\frac{7}{8}$]
Peer–peer communication prob. (P_{p2p})	0.2 / [0.1, 0.5]
Peer–merchant comm. prob. (P_{p2m})	0.6 / [0.1, 0.5]
Peer mobility prob. (P_{move})	0.2 / [0.05, 0.7]
Starting account balance ($AS_{balance}$)	500 / 200
Number of agents ($\|A\|$)	50000 / 1000
Number of merchants (M)	30
Number of banks (B)	5
Online bank contact prob. (P_{bank})	0.01
Merchant sync frequency (F_{sync})	8 steps
Top-up amount (M_{topup})	10 units
Offline transaction limit ($L_{offline}$)	6
Low balance top-up threshold (T_{low})	2 units
Total simulation steps (S_{steps})	72 steps
Number of graph nodes ($\|V\|$)	64

lished channels in the Lightning Network enable peer-to-peer transactions without connectivity to the Bitcoin ledger.

B. Experimental and simulation setups

We conduct two types of experiments: *threat level experiments* and a *sensitivity analysis*. The first evaluates key system metrics, such as counterfeit and double-spender detection against different threat levels (*i.e.*, different percentages of malicious agents in the system). Sensitivity analysis varies the values of the four input parameters of the simulation (*i.e*, peer to peer/merchant interaction probabilities, peer movement probability, and the ratio of honest agents to double-spenders) to understand the relationship between its input and output parameters. The simulation environment here was developed by using a modified version[1] of Briolette [8] codebase. The modifications yielded a significant performance boost over the default Briolette simulator configuration, enabling us to perform large-scale simulations with up to 50 thousand user agents in a reasonable timeframe. Additionally, an improved transaction fork detection algorithm has been added to the simulator since the original algorithm ignored certain double-spending behavior, resulting in skewed measurements.

Table II presents an overview of the parameters used in the experiments. The first part of the table summarizes parameters that vary between the threat level and sensitivity experiment types, with corresponding values separated by a slash character. For sensitivity experiments, the $[a, b]$ notation represents an inclusive range with uniform distribution for parameter values used throughout the evaluation. The second part of Table II presents parameters that are identical across both experiments.

In all of the conducted experiments, all consumer agents start with a predefined amount of coins in their bank account, and each coin has a denomination of one unit of currency. Consumers are allowed to perform $L_{offline}$ transactions without having to reconnect to the bank, and if the user-held balance becomes too low (less than T_{low}), the agent will request a top-up from the bank while also syncing their data with it. Every merchant in the network interacts with a bank every F_{sync} simulation steps (one simulation step is equivalent to one hour) by going online. In the real world, merchants typically have constant access to the network, meaning that they are not restricted to interacting with the bank at periodic intervals, except in certain scenarios such as farmer's market, where connectivity can be limited. Such activity could correspond to depositing profits at the end of the business hours or accounting period depending on the time scale. During this interaction, the merchant shares the data about all coins that have been spent during the window between synchronizations, which the bank later checks for counterfeits. During this interaction, the merchant also synchronizes their epoch data with the bank. Malicious agents enter the simulation in two batches at steps 0 and 15.

Both the threat level and sensitivity analysis experiments are conducted on a network of $|V|$ nodes with M merchants and B banks. The experiments involve $|A|$ agents with $AS_{balance}$ coins of initial balance. The edge configuration of the network is fixed and generated using either the Barabási-Albert or Watts-Strogatz graph generation algorithms. To reduce the noise in the sensitivity analysis results, the simulation is run 10 times for each of the 20,480 unique parameter combinations, and the results are averaged out.

The parameter values in Table II were selected to approximate conditions typically observed in both urban and rural scenarios, though they may not fully capture real-world heterogeneity. Each additional parameter increases the dimension space of the sensitivity analysis, increasing the simulation complexity. Ergo, where direct evidence was unavailable, the values were determined using heuristics. For instance, the values for $Ratio_{ds}$ (ratio of double-spenders to the overall number of agents in the system) in threat level experiments span from a relatively adversarial configuration—one malicious agent per ten honest agents—to a more extreme setting of one malicious agent per one honest agent. The latter represents an adversarial intensity during which honest agents in distributed systems are typically unable to maintain state consensus. In practice, parameter values for deployment could be obtained either directly through empirical studies or indirectly via appropriate proxy metrics.

C. Threat level experiments

Threat level experiments describe the behavior of the system under different ratios of malicious agents in the system (see Table II). Specifically, the experiments examine the ratio of detected counterfeit tokens to the total number of tokens, the percentage of double-spenders detected, and the mean difference between the latest epoch issued by the operator and the epochs of the user agents in the field for each threat level.

1) Discovered counterfeit ratio: Figure 2a presents how well the CBDC system detects and seizes counterfeit coins for the urban and rural scenarios. The discovered counterfeit ratio is calculated as the total value of seized tokens over the initial token supply of the system: $DCR = \frac{\sum_{i=1}^{N} V_i^{\text{cf}}}{TS_{\text{init}}}$, where N is the total number of seized counterfeit coins, V_i^{cf} is the value of seized counterfeit coin i (in our case, each coin is worth 1 unit of currency), and TS_{init} is the initial token supply—that is, the total valuation of coins generated by the simulation before agents begin interacting.

It is observed that all threat levels follow the same pattern: an initial lag period is observed before the first merchant-

[1]https://github.com/ukitta555/briolette_cbdc_paper

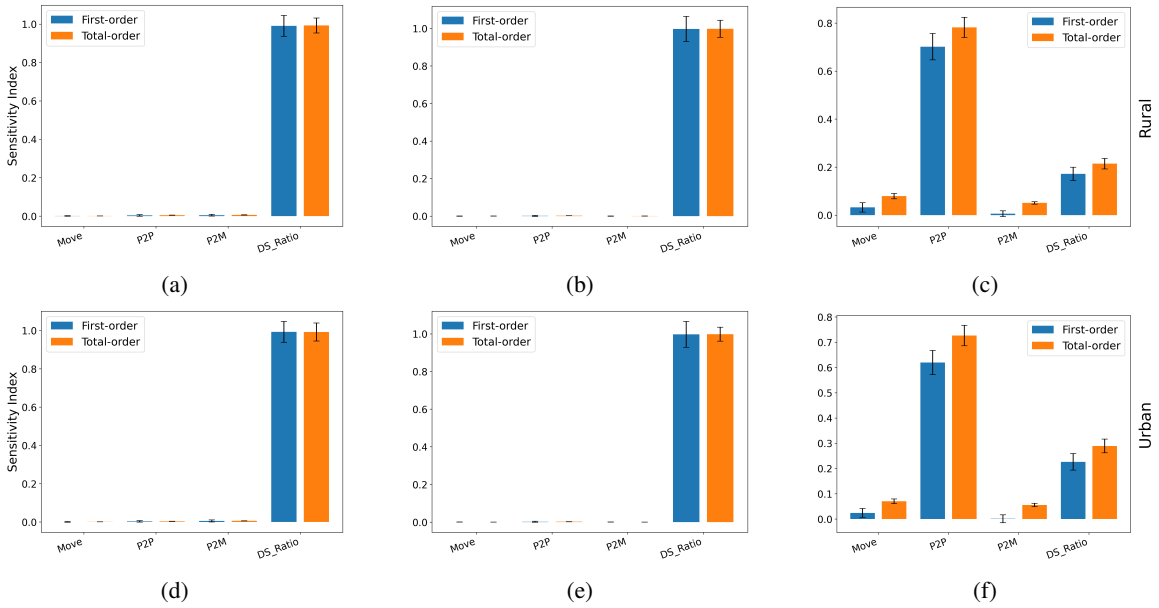

Fig. 4: First-order and total-order Sobol indices for (a,d) DCR, (b,e) DS_{caught} and (c,f) E_{diff}.

to-bank interaction occurs, followed by a linear ascend, and an inflection point after which the counterfeit ratio increases logarithmically until saturation. The inflection point represents the point in time when most of the malicious agents have been detected and revoked by the CBDC system, and the bank only has to collect the double-spent coins that are still circulating in the network from previous double-spend events.

It is also observed that with an increase in threat level, the counterfeit ratio increases as well. In particular, the relationship between threat levels and the counterfeit ratio appears to be linear—a two-fold increase in the ratio of double-spenders to honest agents in the system leads to a roughly two-fold increase in counterfeit detection.

The results for the rural scenario paint a similar picture, with the only major difference being the maximum values for each of the threat levels. Due to the lower average node degree of the WS graph, it is harder for counterfeit coins to reach the bank, leading to an increase in the quantity of detected counterfeits by the end of the simulation.

2) Caught double-spenders: Figure 2b describes the percentage of malicious agents that have been labeled as double-spenders for the urban and rural scenarios, denoted as DS_{caught}. This figure can be used to interpret the ability of the system to detect malicious agents in a timely manner.

All curves resemble a step-like function, with constant segments and periodical discontinuities. The discontinuities can be explained by the periodic nature of merchant-to-bank communication, when the merchants validate their coins with the ledger, which, in turn, marks the agents that double-spent as malicious.

The discontinuities increase sharply later in the simulation for all curves. This is due to the fact that the second batch of double-spending agents appears in the system on step 15, and there is a lag before they start spending counterfeit tokens. Once those coins are in circulation, the banks will start flagging both the initial batch of double-spenders (that have appeared at step 0) and the batch that has appeared at step 15. This effect is more pronounced for threat levels $\frac{1}{5}$, $\frac{1}{3}$, and $\frac{1}{2}$. Also note that by the end of the simulation, the discontinuities stop appearing. This is because all of the double-spenders have been caught—for all lines in Figure 2b, the final ratio value is exactly the percentage of the malicious agents in the system.

It is worth noting that, in a real-world deployment, additional controls - other than online checks - would be in place to prevent and/or detect double-spending (*e.g.,* secure hardware, reputation checks), reducing the time needed to detect a double-spender.

3) Epoch differences: Figure 3 presents the mean and standard deviation of the difference between the latest epoch number published by the system and the epoch number known to a random honest agent, which is denoted as E_{diff}. Here, one can reason about the propagation of information in the network as follows: bigger/smaller values of the mean represent slower/faster information propagation, while the standard deviation represents how far apart the agent population is from a random agent.

Four curves, split into two pairs, are depicted in each of the sub-figures. The blue/green pair represents the unnormalized mean and reflects the absolute epoch difference. Similarly, the orange/red pair depicts the standard deviation of the difference. Both pairs have a periodic structure related to the merchant-to-bank synchronization cycles. The range of values depicted in the curves is shown to be stable throughout the simulation, implying that the information disparity is eliminated at the end of each period (and does not grow over time).

The maximum and minimum values of the mean span the interval from 0 to 1, indicating the network achieves consensus at the latest epoch number. The spikes in the mean curves are caused by epoch updates arriving from the banks to merchants. At the start of the period, the mean difference is 1 indicating maximum disparity between the bank and network nodes. As time passes, the information disseminates across the network,

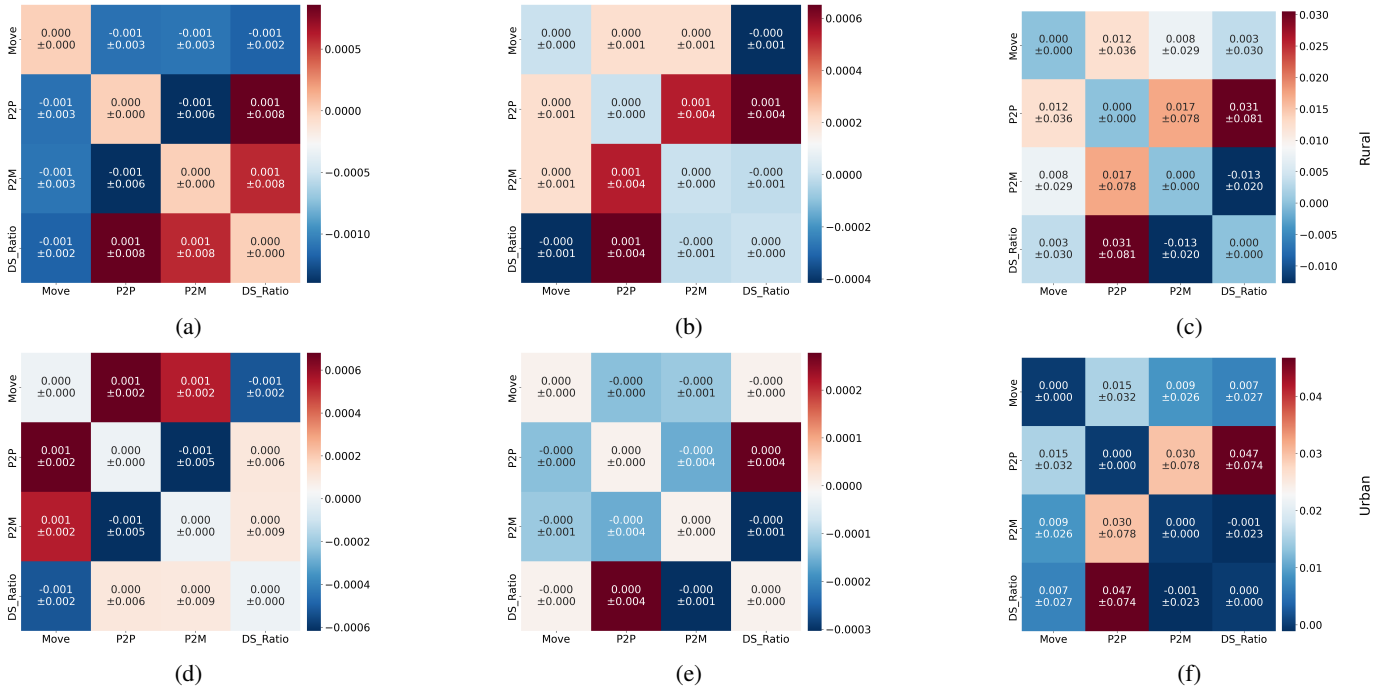

Fig. 5: Second-order Sobol Indices for (a,d) DCR, (b,e) DS_{caught} and (c,f) E_{diff}.

as shown by the decreasing mean difference values.

The standard deviation curves also support the hypothesis that the system is stable in terms of information propagation. A delay can be observed between the peak of the mean and standard deviation curves that is attributed to the propagation between the bank update and information flow through the network. In essence, the peak of the standard deviation curve represents the point in time where the network is roughly bifurcated, where 50% of the nodes have received the update and the other 50% have not. It can be observed that a percentage of nodes always maintains a disparity with the bank's information throughout the simulation's life-cycle.

Rising threat levels lead to higher minimums for both the mean and standard deviation in each period. This is because malicious agents withhold information from honest agents, always sharing information about the first epoch when transacting. This makes the information propagate slower across the network, which causes the mean and standard deviation to increase. However, note that minimums trend down as the simulation continues, which confirms our hypothesis: as the number of malicious agents in the system reduces, the information spreads more freely across the network. This is evident in Figure 3a, where the curve goes to zero before the simulation is over since all double-spenders have been caught.

D. Sensitivity analysis experiments

This section describes the sensitivity analysis results. Figure 4 presents the results for Sobol indices analysis [26], which is a standard way to estimate the impact of simulation inputs on an output variable using variance decomposition. Each subplot consists of four pairs of bar segments, with each pair representing one of the four inputs that we perform the analysis on. The inputs are peer to peer/merchant interaction probabilities, peer movement probability, and the ratio of malicious agents (see Table II). The inputs are located on the horizontal axis. The analysis is performed for three outputs, namely discovered counterfeit ratio, double-spenders caught ratio, and mean epoch difference, presented previously in threat level experiments. The vertical axis represents the sensitivity values. For example, the blue bars are the first-order Sobol indices, which represent the effect on the output explained specifically by the input to which the bar is related. However, this metric needs to be evaluated jointly with another—the total Sobol indices, which are represented by orange bars. The total Sobol indices show the effect of all the groups of variables that contain the input in question on the output. If the first-order index is of the same magnitude as the total index, then the input contributes a lot to the variance of the output. Otherwise, a higher-order analysis is typically performed to take into account all pairwise interactions between inputs.

For both urban and rural scenarios, the inputs that make the most impact are the same: it is the ratio of double-spender agents for the counterfeit and double-spenders caught ratios, and the global to local mean epoch difference is best explained by peer to peer interaction probability. In all three instances, the first-order indices are very close to total-order indices, indicating that there are no meaningful higher-order interactions happening with those parameters for that output. This is confirmed by second-order Sobol indices in Figure 5, where all of the input interactions failed to produce an effect larger than 0.05 and most of the interactions had no effect on the output in all three cases for both scenarios.

Scatter plots for the counterfeit, epoch difference, and double-spenders outputs are presented as evidence to bolster support for the initial hypothesis. Each plot depicts direct relationship with the ratio of malicious agents or peer-to-peer

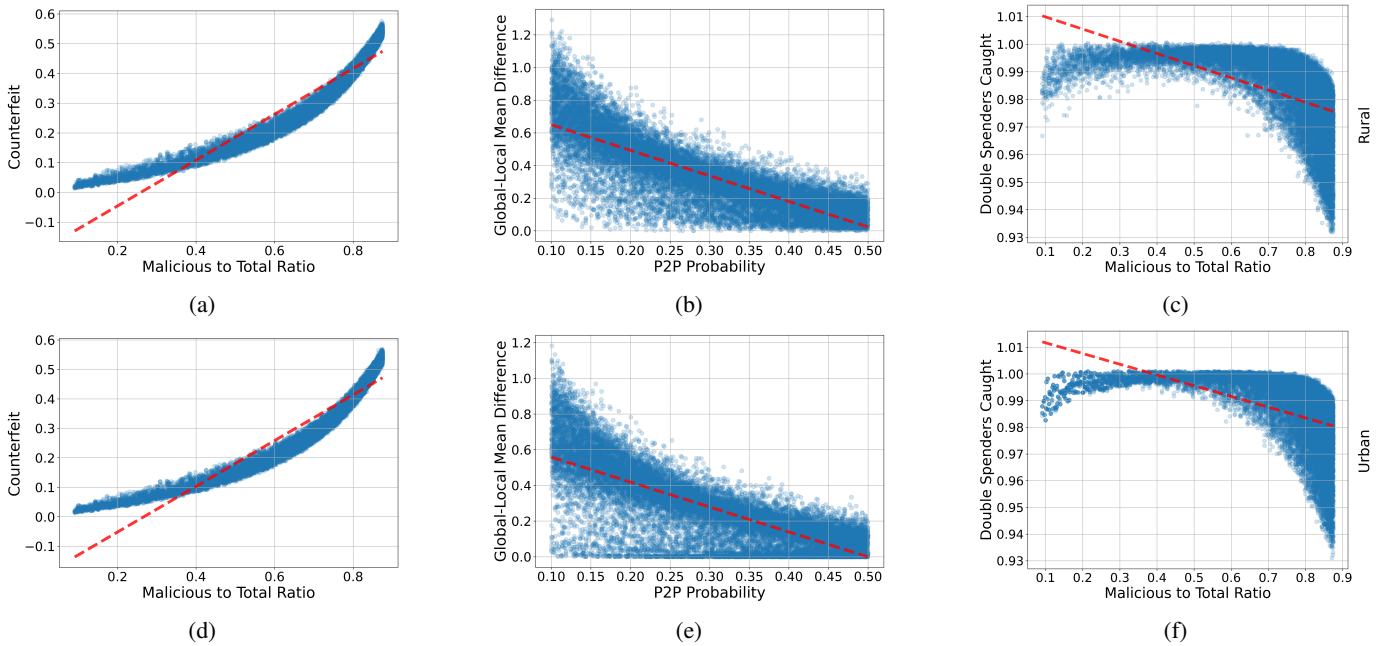

Fig. 6: Scatter plots for (a,d) DCR vs. $Ratio_{ds}$, (b,e) E_{diff} vs. P_{p2p}, and (c,f) DS_{caught} vs. $Ratio_{ds}$.

interaction probability. Figures 6a, 6d show the relationship between the counterfeit discovered and the ratio of malicious agents, where the latter parameter ranges from $\frac{1}{11}$ to $\frac{7}{8}$. The result shows a strong positive correlation for both urban and rural scenarios, which validates the hypothesis that there is a considerable impact on the output variable. The result is expected since increasing the number of malicious agents in the tested system increases the amount of double-spent coins created in each step of the simulation. As a side note, the growth is not entirely linear—it speeds up closer to more extreme values of the malicious agents ratio.

Figures 6b, 6e present a negative correlation between peer-to-peer interaction probability and the mean difference between the global epoch and the local epoch, where the interaction probability ranges from 0.1 to 0.5. The plots for the rural and urban scenarios share a lot of similarities. For example, both Figure 6b and 6e show a large variance for the results corresponding to small values of the interaction probability. The variance greatly decreases as the interaction probability grows, which aligns with the hypothesis that more interaction between honest agents results in faster data propagation across the network. However, differences across figures are also present, with the main one being the numerical values of the epoch difference. The rural scenario shows a bigger difference compared to the urban one, which can be explained by differences in graph structure.

Finally, Figures 6c & 6f present the relationship between the ratio of discovered double-spenders and the ratio of malicious agents in the system, where the input still ranges from $\frac{1}{11}$ to $\frac{7}{8}$. As in the case with global and local epoch difference, the figures show strong negative correlation, which validates the hypothesis about the impact of the input on the output. The sign of the correlation can be interpreted in an intuitive manner since the more malicious agents there are in the system, the harder it gets to catch the double-spenders as they start to double-spend with other malicious agents, who will never spend the coins of other malicious agents. The outliers in the tail range from $\frac{1}{11}$ to $\frac{2}{5}$ can be explained by the specifics of the simulation environment—some of the malicious agents do not get the chance to double-spend at all throughout the simulation if they do not move to a location with an agent. As the number of malicious agents grows, this uncaught minority represents a smaller percentage of malicious agents, hence an increasing trend is present. However, at the inflection point, more and more malicious agents who have interacted with the honest ones do not get caught, and the trend reverses.

V. Conclusion

Offline functionality in digital currency systems is a highly desired feature that promotes resilience and improves financial accessibility for unbanked communities, but also introduces a variety of operational risks. This paper presented a flexible and comprehensive framework for the risk assessment of offline digital currency systems that can be used as a tool for making informed policy decisions. Through the use of a wide range of configurable parameters, including the composition of the agent population, the probabilities of peer interactions, and the synchronization intervals, the framework allows a detailed evaluation of various digital currency designs. By applying the proposed framework to an open source CBDC prototype, we demonstrated its ability to provide insight on how different parameter configurations affect key performance metrics, such as counterfeit detection rates, double-spending incidents, and the propagation of state updates.

Future work in this domain may include formal methods and analytical modeling from a risk quantification perspective to gain a better understanding of the trade-offs between the impact of threats and the costs to prevent them.

References

[1] DefiLlama. Accessed: Jun. 01, 2025. [Online]. Available: https://defillama.com/

[2] VettaFi. Cryptocurrency ETF list. Accessed: Jun. 01, 2025. [Online]. Available: https://etfdb.com/etfs/currency/cryptocurrency/

[3] "A proposal for a retail central bank digital currency (CBDC) architecture," Bank for International Settlements, Tech. Rep., Dec. 2024. [Online]. Available: https://www.bis.org/publ/othp89.htm

[4] "Central banks and distributed ledger technology: How are central banks exploring blockchain today?" World Economic Forum, Tech. Rep., Mar. 2019. [Online]. Available: https://www3.weforum.org/docs/WEF_Central_Bank_Activity_in_Blockchain_DLT.pdf

[5] "Central bank digital currency (CBDC): Retail considerations," Payments Canada, Tech. Rep., Aug. 2022. [Online]. Available: https://www.payments.ca/sites/default/files/2022-08/PaymentsCanada_2021CBDCRetailConsiderations_En.pdf

[6] "Project Polaris: Handbook for offline payments with CBDC," Bank for International Settlements, Tech. Rep., May 2023. [Online]. Available: https://www.bis.org/publ/othp64.htm

[7] "Central bank digital currencies: Foundational principles and core features," Bank for International Settlements, Tech. Rep., Oct. 2020. [Online]. Available: https://www.bis.org/publ/othp33.htm

[8] W. Drewry, "GitHub - google/briolette: Briolette is an experimental framework for researching offline digital currency designs," github.com/google/briolette, 2023, accessed: May 06, 2025.

[9] S. Allen, S. Čapkun, I. Eyal, G. Fanti, B. Ford, J. Grimmelmann, A. Juels, K. Kostiainen, S. Meiklejohn, A. Miller, E. Prasad, K. Wüst, and F. Zhang, "Design choices for central bank digital currency: Policy and technical considerations," National Bureau Of Economic Research, Tech. Rep., Aug. 2020.

[10] C. Barontini and H. Holden, "Proceeding with caution - a survey on central bank digital currency," Bank for International Settlements, Tech. Rep., Jan. 2019. [Online]. Available: https://papers.ssrn.com/abstract=3331590

[11] U. Bindseil, "Tiered CBDC and the financial system," European Central Bank (ECB), Tech. Rep., 2020. [Online]. Available: https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2351~c8c18bbd60.en.pdf

[12] R. Auer, G. Cornelli, and J. Frost, "Rise of the central bank digital currencies: drivers, approaches and technologies," Aug. 2020. [Online]. Available: www.bis.org

[13] A. Carstens, "Digital Currencies and the Future Monetary System," *Hoover Institution Policy Seminar*, vol. 89, no. 1, p. 17, 2021. [Online]. Available: https://www.bis.org/speeches/sp210127.pdf

[14] J. Gross, J. Sedlmeir, M. Babel, A. Bechtel, and B. Schellinger, "Designing a central bank digital currency with support for cash-like privacy," 2021. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3891121

[15] D. A. Zetzsche, R. P. Buckley, and D. W. Arner, "Regulating libra," *Oxford Journal of Legal Studies*, vol. 41, no. 1, pp. 80–113, Dec. 2020.

[16] H. Armelius, C. A. Claussen, and I. Hull, "On the possibility of a cash-like CBDC," 2021. [Online]. Available: https://ideas.repec.org//p/zbw/esprep/231485.html

[17] Y. Chu, J. Lee, S. Kim, H. Kim, Y. Yoon, and H. Chung, "Review of offline payment function of CBDC considering security requirements," *Applied Sciences*, vol. 12, no. 9, p. 4488, Apr. 2022.

[18] T. Alper, "Further details of 'offline' Chinese Digital Yuan 'hard wallet' emerge," 2021. [Online]. Available: https://cryptonews.com/news/further-details-of-offline-chinese-digital-yuan-hard-wallet-8891.htm

[19] A. Tsareva and M. Komarov, "Retail central bank digital currency design choices: Guide for policymakers," *IEEE Access*, vol. 12, pp. 66 129–66 146, 2024.

[20] A. Ramadiah, M. Galbiati, and K. Soramäki, "Agent-based simulation of central bank digital currencies," *SSRN Electronic Journal*, 2021.

[21] C. León, J. F. Moreno, and K. Soramäki, "Simulating the adoption of a retail CBDC," *Jahrbücher für Nationalökonomie und Statistik*, Jul. 2024.

[22] S. B. Espedal and D. A. Janzso, "Design choices for offline transactions in a norwegian central bank digital currency," Master's thesis, Norwegian University of Science and Technology, 2022. [Online]. Available: https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/3041390

[23] W. Bian, Y. Ji, and P. Wang, "The crowding-out effect of central bank digital currencies: A simple and generalizable payment portfolio model," *Finance Research Letters*, vol. 43, p. 102010, Nov. 2021.

[24] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, Oct. 1999.

[25] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, Jun. 1998.

[26] I. M. Sobol, "Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates," *Mathematics and computers in simulation*, vol. 55, no. 1-3, pp. 271–280, 2001.