

VILNIUS UNIVERSITY
FACULTY OF ECONOMICS AND BUSINESS ADMINISTRATION

STRATEGIC ECONOMICS

Jonas Mekšraitis

MASTER THESIS

**VIETINIŲ VALIUTŲ
EKONOMINIO RACIONALUMO
PAIEŠKOS: MOKYKLINĖS
VALIUTOS INICIATYVOS
ATVEJO ANALIZĖ**

**IN SEARCH OF THE RATIONALE
FOR LOCAL CURRENCIES: A
CASE STUDY OF A SCHOOL
CURRENCY INITIATIVE**

Supervisor Prof. Dr. Robertas Zubrickas

Vilnius 2026

TABLE OF CONTENTS

INTRODUCTION	4
1. PARALLEL CURRENCIES: CONCEPTS, WHEN THEY MATTER, AND WHY BARTER DOESN'T HAPPEN.....	6
1.1 Parallel currencies: concepts and the basic design problem.....	6
1.1.1 The problem: why barter usually fails to organize exchange	7
1.1.2 When parallel currencies matter	8
1.1.3 Design levers and archetypes of parallel currencies.....	9
1.1.4 Coordinating Exchange in Small Networks: Tokens over Barter	11
1.2. Small-community settings and the leakage problem	12
1.2.1 Mechanism of leakage	12
1.2.2 The “forced-barter” intuition and its limits	13
1.3 Utility of money per se.....	14
1.4 From micro to macro: how MPR eats MPC	16
1.4.1 A minimal accounting frame	16
1.4.2 What redeemability does in practice (evidence by archetype)	17
1.4.3 From levers to outcomes: linking mechanisms to behavior	18
1.4.4 Measuring and diagnosing token performance.....	19
1.4.5 Summary:.....	20
2. CONCEPTUAL FRAMEWORK.....	21
2.1 Introduction.....	21
2.1.1 Preliminaries and notation	22
2.1.2 Edgeworth efficiency with a separable money good.....	23
2.1.3 Distributional effects in competitive equilibrium (budgets).....	24
2.1.4 Symmetric benchmark (closed form)	25
2.1.5 General ceteris paribus (unequal incomes and/or real-goods tastes).....	26
2.2 A Market-Demand Proof: Utility-Bearing Money Reduces Real-Goods Demand.....	28
2.3 General-Equilibrium Clarifications	29
2.4 Non-Hoardable Currency (Demurrage) and the Reversal.....	30
2.5 Conclusion	30
3. RESEARCH DESIGN AND EMPIRICAL STRATEGY	31
3.1 Research design.....	31
3.2 Setting and participants:	32
3.3 Data collection methods:	33
3.4 Intervention Design.....	33
3.4.1 Collectible Cards	34
3.4.2 Active Buying Phase (Non-Structural).....	34
3.4.3 Ethical Boundaries.....	34
3.5 Data Analysis Strategy	35
3.5.1 Dataset structure and preprocessing	35
3.5.2 Descriptive market outcomes	35
3.5.3 Currency flows and low-velocity proxies.....	35
3.5.4 Pre/post intervention comparison (two-phase accounting).....	36
3.5.5 Qualitative observations	36
3.6 Limitations.....	36

4. DESCRIPTIVE FINDINGS FROM THE PILOT	38
4.1 Overview of the experiment and dataset	38
4.2 Market composition and participation	38
4.3 Trade volume, sales rates, and time to sale	40
4.4 Currency flows and low velocity in the real-goods segment	41
4.5 Intervention results: collectible cards as a spending nudge.....	42
4.6 Heterogeneity: age, gender, and item type	43
4.6.1 Age and trading outcomes for real goods	43
4.6.2 Gender and trading outcomes for real goods	44
4.6.3 Inventory composition by gender: books vs. non-books	44
4.6.4 Inventory composition by age: books vs. non-books	45
4.6.5 Linking real-goods heterogeneity to the intervention.....	46
4.7 Synthesis of empirical findings	46
5. IMPLICATIONS, LIMITATIONS, AND FUTURE RESEARCH	48
5.1. Summary of key findings	48
5.2 Parallel currencies, leakage, and acceptance density	49
5.3 Hoardability and the behavioral value of waiting	50
5.4 Market thinness, inventory composition and distributional patterns	51
5.5 Implications for local-currency design and for school-based experiments	52
5.6 Limitations and avenues for future research	53
REFERENCES:	55
SUMMARY	59
SANTRAUKA	60
APPENDIXES:.....	62

INTRODUCTION

Local and complementary currencies are typically introduced when exchange is possible in theory but weak in practice. In small communities, trade can stay thin because goods are not consistently available, and holding money is often more attractive than spending it. A closed-loop token system is meant to lower coordination costs and make exchange feel “worth it,” especially for low-value goods that people would not bother selling for euros. This thesis examines that logic in a school setting by implementing a non-redeemable currency and tracking what happens to circulation, listings, and participation over time.

Research on parallel currencies covers local civic currencies, mutual-credit systems, and other closed exchange arrangements. The main lesson across cases is consistent: design details determine whether a currency becomes a medium of exchange or turns into something people simply hold. Acceptance density, redeemability, and the availability of things worth buying matter most. Much of the evidence remains case-based and difficult to generalize, especially in small markets where outcomes are heavily shaped by inventory variety and participation patterns. This thesis contributes by observing those micro-dynamics directly through transaction records in a contained environment.

The research problem addressed is that a non-convertible local currency may enable exchange while still drifting into low circulation if users prefer to hoard tokens or if supply remains narrow. The research question is: can a school-based, non-redeemable currency increase exchange activity in a thin market, and which constraints most strongly limit circulation in practice? The working expectation is that tokenization can make certain low-value goods tradable by reducing friction, but that circulation will remain limited if attractive inventory is scarce and holding tokens has high option value.

The aim of the thesis is to assess whether a school currency can mobilize idle goods for exchange and to identify the mechanisms shaping circulation and participation. The objectives are:

- 1) summarize key design levers discussed in the parallel currency literature
- 2) frame hoarding and low velocity using a simple conceptual benchmark
- 3) implement an embedded pilot in a school setting with clear rules and ethical boundaries

4) analyze an item-level dataset to describe listings, sales performance, participation concentration, and the effect of a spending nudge introduced during the pilot.

Methodologically, thesis combines literature analysis, a conceptual/formal framing of hoarding incentives, and a descriptive case study based on transaction logs. The pilot ran for eight weeks (excluding one holiday week) in a private school in Vilnius with a potential pool of approximately 115 students (grades 5–11). Transactions were recorded in a shared spreadsheet as an item-level log of goods entering and leaving the shop; after cleaning, the analytical sample contains 64 entries. During the pilot, a low-cost intervention was introduced to address stagnation: collectible cards sold for tokens, designed to create an additional spending outlet without importing external material value.

The thesis is structured in five chapters:

- 1) reviews the literature of parallel currencies and circulation constraints.
- 2) develops the conceptual benchmark for hoarding and low velocity.
- 3) describes the research design, setting, data, and ethical constraints.
- 4) reports descriptive results and compares pre- and post-intervention dynamics.
- 5) summarizes implications, limitations, and directions for future research.

The main limitations arose from participation being voluntary and uneven, resulting in modest dataset, and inventory composition changed over time in ways that likely affected demand and willingness to list new goods. Pricing was set pragmatically rather than through a standardized valuation method. Given that participants were minors, anonymity was maintained and no personal data were collected, which prevents individual-level tracking of balances and repeated behavior. Replication is constrained by access, as the pilot relied on the researcher's institutional position.

AI tools were used as writing support. GPT-based models were used to improve clarity, grammar, and academic phrasing, and to propose alternative wording and paragraph structures. The author developed the ideas and argument structure, selected and edited the final phrasing, and verified the content.

1. PARALLEL CURRENCIES: CONCEPTS, WHEN THEY MATTER, AND WHY BARTER DOESN'T HAPPEN

Complementary and local currency initiatives are often presented as practical responses to a recurring constraint: communities may have willing buyers and sellers, yet exchange remains thin when liquidity is scarce, purchasing power leaks outward, or coordination costs are high. Because the relevant research spans case studies, monetary theory, and policy evaluations, it can easily read like a catalogue. Given the context main guiding question is: which design features determine whether a token system circulates persistently within a limited acceptance environment, rather than stalling, being redeemed, or failing to localize value?

Firstly, why tokens can dominate informal alternatives such as barter in decentralized settings? Monetary search models emphasize that barter requires a double coincidence of wants and repeated matching effort; even intrinsically worthless units can improve outcomes by lowering search and timing frictions (Kiyotaki and Wright, 1993; Lagos and Wright, 2005). Secondly, it synthesizes evidence across the main design archetypes: mutual credit, par-redeemable tokens, and voucher-like instruments, highlighting the trade-off between lowering adoption frictions and preserving on-network circulation, with particular attention to acceptance density and the extent of spendable goods (Stodder and Lietaer, 2016; Fleischman and Dini, 2020; Kim and Oh, 2021; Hua, 2022). Third, it translates these designs into a compact mechanism frame - circulation is governed by a small set of margins: opportunity to spend, incentives to hold and, where relevant, incentives to exit, alongside leakage as the loss of value from the intended loop.

Lastly, how the applied literature typically diagnoses token performance: through acceptance redemption, holding dynamics, and leakage? So that later empirical analysis can be interpreted as evidence about mechanisms rather than as isolated case comparisons.

1.1 Parallel currencies: concepts and the basic design problem

Parallel (or “complementary”) currencies are monetary instruments that circulate alongside national currency but are intentionally restricted by acceptance network, geography, redeemability, or time to serve goals such as liquidity provision, stabilization, or localization (Blanc, 2011; Seyfang and Longhurst, 2013). One question that matters for the rest of the thesis: what makes a token system

circulate inside a limited acceptance network, rather than stalling, being redeemed, or failing to localize value?

A useful starting point is the counterfactual. In small or liquidity-constrained settings, the realistic alternative is rarely a smooth return to barter. More often, exchange becomes thinner and more delayed unless some commonly accepted medium of exchange or credit arrangement emerges. The sections below move from that baseline to the main currency archetypes and then to the design levers that repeatedly determine whether circulation becomes self-sustaining.

1.1.1 The problem: why barter usually fails to organize exchange

It is sometimes assumed that in small communities, direct barter could replace money once national currency leaks out of the local circuit. Classic monetary theory explains why that expectation rarely holds. In decentralized exchange, barter depends on a double coincidence of wants (I must have what you want *and* want what you have) as well as time to search, verify quality, and enforce informal contracts. Search-theoretic models formalize these frictions, showing that even intrinsically worthless tokens can outperform barter by reducing coordination and matching costs (Kiyotaki and Wright, 1993). Later work embeds the same logic in a tractable macro framework where decentralized meetings alternate with centralized markets (Lagos and Wright, 2005). Both approaches predict that as trading frictions increase, the welfare gains from a commonly accepted medium of exchange rise disproportionately.

Empirical work broadly confirms this intuition. In laboratory and field settings calibrated to monetary search models, participants trade more frequently and reach higher welfare under monetary exchange than under gift- or barter-based environment, even when non-monetary allocations are theoretically feasible (Duffy and Puzzello, 2014; Rötheli, 2009). In practice, barter does appear, but briefly. It tends to collapse once tokens emerge that provide liquidity services.

When national currency drains from a small economy, the realistic counterfactual is not that people will “just barter,” but that exchange will contract or shift toward informal credit unless a simple medium of exchange is available. This is where parallel or community currencies find their rationale: they lower matching costs without depending on fragile barter coincidence of wants (Kiyotaki and Wright, 1993; Lagos and Wright, 2005)

1.1.2 When parallel currencies matter

Parallel (or “complementary”) currencies become relevant when a community faces liquidity shortages, credit rationing, thin markets with few counterparties, or policy goals that require spending to stay within a defined network or flow toward selected merchants. Evidence spans several distinct families of systems as discussed below.

One influential archetype is business-to-business mutual credit, where firms create liquidity internally through multilateral trade credit rather than through redeemable tokens. The Swiss WIR, established in the 1930s, and more recent networks such as Sardex in Italy, extend trade credit multilaterally within closed loops. Empirical studies show that WIR activity is countercyclical and helps stabilize relationships among small and medium-sized firms. Recent research on Sardex finds that obligation clearing and mutual credit reduce internal net debt and strengthen firms’ resilience to liquidity shocks - outcomes consistent with theoretical predictions that a parallel unit can ease cash constraints without enabling leakage (Stodder and Lietaer, 2016; Fleischman, 2020; Reyns et al., 2025).

A second archetype is public voucher-style instruments and municipal e-currencies, designed to stimulate local spending through redeemable, often time-limited units tied to specific merchants or sectors. During and after the COVID-19 pandemic, some governments introduced redeemable local vouchers, often time-limited and restricted to specific merchants. Studies from Taiwan and South Korea find mixed but generally positive short-term effects, with spending increases concentrated in targeted sectors and months. These programs highlight a recurring design tension: redeemability improves usability but also accelerates leakage back into national currency unless merchant restrictions are binding (Hua, 2022; Kim and Oh, 2021).

A third archetype consists of community currencies explicitly aimed at “localization,” using 1:1 redeemable tokens and soft incentives to encourage residents to keep spending within the local area. Initiatives in the United Kingdom and the United States, such as the Bristol Pound and BerkShares, attempted to retain value locally through 1:1 redeemable tokens coupled with soft incentives. Rigorous evaluations are cautious: for BerkShares, a synthetic-control study detects no measurable macroeconomic impact, while for the Bristol Pound, both academic and practitioner analyses identify institutional and adoption barriers that limited reach (Matti and Zhou, 2022;

Marshall, 2018). These cases suggest that without broad acceptance and strong incentives, redeemable local currencies struggle to shift procurement patterns or supplier networks.

In short, with few exceptions, parallel currencies have meaningful effects when they either generate internal liquidity that cannot easily exit (as in mutual-credit systems) or constrain spending through time, merchant, or geographic rules (as with vouchers). When neither condition is met, their impact tends to dissipate quickly.

1.1.3 Design levers and archetypes of parallel currencies

“Parallel currency” serves as an umbrella term for monetary instruments that circulate alongside national currency but are intentionally restricted: by acceptance network, geography, redeemability, or time, to achieve specific goals such as liquidity provision, stabilization, or localization. The literature variably labels these instruments as complementary, community, or local currencies; here “parallel” is used to emphasize their coexistence with, rather than opposition to, the national unit (Blanc, 2011; Seyfang and Longhurst, 2013). Parallel currencies can be classified along several main design axes.

Unit and redeemability. Many local voucher systems and city currencies are redeemable at par with the national unit. By contrast, mutual credit systems, whether community-based local exchange trading systems (LETS) or business networks such as WIR and Sardex, are not redeemable into fiat. WIR and Sardex generally maintain a 1:1 accounting peg to the national unit (CHF or EUR), whereas many LETS and time-based currencies, such as Ithaca HOURS, denominate balances in hours or service credits, functioning as reciprocal accounting systems rather than cash-redeemable money (Blanc, 2011; Seyfang and Longhurst, 2013; Stodder and Lietaer, 2016; Fleischman and Dini, 2020).

Redeemable local currencies (par tokens): The Bristol Pound (UK, 2012–2021) operated in both paper and digital form, allowed limited local tax payments, and eventually wound down under adoption, cost, and digital-competition pressures. Evaluations describe high visibility, but limited success in redirecting procurement toward local suppliers (Marshall, 2018). BerkShares (US, since 2006) similarly offer par redeemability with a purchase discount, yet a synthetic-control study finds no macroeconomic effects on local business dynamics or county indicators, suggesting that modest price incentives, absent supply-chain restructuring, are insufficient to shift spending patterns at scale (Matti and Zhou, 2022).

Non-redeemable mutual credit (network-redeemable): The Swiss WIR is a long-standing business-to-business clearing network operating in a parallel unit. Empirical evidence points to macro-stability and countercyclical behavior among participating SMEs (Stodder and Lietaer, 2016). Sardex (Italy) offers a similar multilateral clearing mechanism; transaction-level data show that obligation-clearing and mutual credit reduce internal net debt and strengthen liquidity, with newer work indicating resilience gains at the firm level (Fleischman, 2020; Reyns et al., 2025).

Social and household systems: Time banks and LETS operate with non-redeemable units (typically hours) that monetize trade and voluntary exchanges. Reviews consistently report social and well-being benefits, though these systems are not designed to influence goods-market procurement or local leakage (Baines et al., 2020).

Holding costs and spending incentives: Spending incentives can be shaped either by penalizing holding or rewarding immediate expenditure. Demurrage or expiry introduces a negative return on balances, such as the Chiemgauer’s historical stamp fee, while bonuses and point-of-sale discounts (often in the 5–10% range, including merchant-funded rebates) encourage rapid circulation and raise on-network velocity (Gelleri, 2009; Hayashi, 2021; Kim and Oh, 2021; Hua, 2022). It employs demurrage, originally through physical stamps, now digitally via time decay, to promote faster recirculation and local multiplier effects. Evidence remains largely case-based, with results dependent on merchant density and the ease of re-spending within the system (Gelleri, 2009; Hayashi, 2021).

Discount and bonus vouchers: Municipal programs such as South Korea’s regional gift certificates and Taiwan’s Triple Stimulus Vouchers show identifiable, short-run spending multipliers concentrated in areas with high acceptance and visible discounts. However, these effects fade quickly without complementary supplier-side or procurement policies (Kim and Oh, 2021; Hua, 2022).

Governance and scope. Governance structures vary by sponsor and objective. Civic or third-sector initiatives (e.g., Bristol, BerkShares, many time banks) are typically run by nonprofits or community groups aiming to foster localization and reciprocity, often with light underwriting and volunteer governance. Business consortia (e.g., Sardex) organize mutual-credit networks among SMEs, offering formal credit lines, curated marketplaces, and active risk management to provide working-capital liquidity. Public or municipal programs (e.g., Korea’s regional gift certificates, Taiwan’s stimulus vouchers) are issued or backed by local authorities as policy instruments. These

are usually geo-fenced or merchant-restricted, time-limited, and managed within administrative budgets and enforcement frameworks (Seyfang and Longhurst, 2013; Marshall, 2018; Matti and Zhou, 2022; Kim and Oh, 2021; Hua, 2022).

Interactions and importance of design levers. Redeemability lowers adoption barriers but raises the share of balances that quickly convert back into national currency which can weaken localization effects. Demurrage or expiry reduces the value of holding but works only when acceptance is dense and supply predictable; otherwise, participants may hoard in anticipation of better future trade opportunities, undermining the “spend-now” effect. Professional governance, as in business groups, can maintain liquidity without fiat leakage, while public voucher programs can stimulate short-term demand but require strong constraints or supplier-side coordination to sustain effects (Gelleri, 2009; Stodder and Lietaer, 2016; Marshall, 2018; Kim and Oh, 2021; Hua, 2022; Matti and Zhou, 2022; Fleischman and Dini, 2020)

1.1.4 Coordinating Exchange in Small Networks: Tokens over Barter

When a community faces a liquidity drought, the practical alternative to a parallel currency is rarely a clean return to barter, but rather thinner and more delayed exchange due to search and timing frictions. Classic monetary search models emphasize that barter requires a double coincidence of wants and repeated matching effort; once these costs are acknowledged, even intrinsically worthless tokens can outperform barter by reducing search time and making exchange more predictable (Kiyotaki and Wright, 1993; Lagos and Wright, 2005). Laboratory evidence is consistent with this mechanism: monetary exchange tends to sustain higher trade volumes and welfare than barter- or gift-like regimes in environments designed to mimic decentralized matching frictions (Duffy and Puzzello, 2014).

This logic also highlights why outcomes in “thin” settings depend heavily on acceptance density and the breadth of the spending menu. Mutual-credit architectures create internal purchasing power without a direct cash-out option, so value must circulate through purchases from other participants. In larger networks such as WIR and Sardex, this structure has been linked to stabilization and relationship-preserving trade during liquidity stress (Stodder and Lietaer, 2016; Fleischman and Dini, 2020). In smaller networks, however, the same design becomes more liable: if few outlets accept the unit or the assortment is narrow, balances can accumulate, and exchange remains sporadic.

Redeemable par tokens and discount vouchers take an opposite approach. Convertibility and simple pricing can reduce adoption frictions, but they also strengthen the incentive to redeem back into national money, making persistence harder unless restrictions are binding (Kim and Oh, 2021; Hua, 2022). The voucher literature illustrates this trade-off: short-run spending effects can be strong, yet the instrument’s durability is limited when it is designed to quickly re-enter standard payment channels. More broadly, par-redeemability can support uptake and reassurance, while simultaneously creating an “exit option” that works against localization unless constraints are sufficiently tight.

Taken together, the literature suggests moving beyond categorical labels (“barter vs. tokens,” “redeemable vs. non-redeemable”) toward the margins that govern circulation: the propensity to transact within the network, the propensity to hold balances intertemporally, and where convertibility exists, the propensity to redeem into national money. Framing designs around these margins provides a compact way to connect micro-level frictions (matching costs, acceptance density, redemption convenience) to macro-level outcomes (velocity, leakage, and the durability of trade).

1.2. Small-community settings and the leakage problem

Small economies rarely struggle because people are unwilling to trade; they struggle because purchasing power fails to stay put. In thin markets, households and small merchants face limited assortments, dependence on external supply chains, and payment systems that make it easy for value to flow back into the national circuit. The result is leakage: income earned locally is not re-spent locally, shortening multiplier rounds and leaving networks shallow.

1.2.1 Mechanism of leakage

Leakage occurs when locally generated income fails to recirculate and instead exits to outside firms, suppliers, or jurisdictions. In regional economics and impact analysis, such leakages are modeled as outflows through imports, non-local profits, taxes paid elsewhere, or savings placed in external institutions. Higher leakage rates reduce the local spending multiplier and dampen subsequent rounds of demand.

Input–output (I–O) frameworks make this relationship explicit - larger the import shares embedded in final or intermediate demand, the smaller the region’s internal linkages and multipliers (IMPLAN, 2023). In applied work, evaluators routinely treat non-local purchases and financial outflows as leakages when estimating local economic impacts (Hughes, 2018; LM3/NEF).

Firstly, assortment and variety are thinner. Under models of monopolistic competition and economic geography, larger markets sustain more product varieties because fixed costs are spread across wider demand; consumers in larger cities therefore face richer menus of goods and services (Krugman, 1991; Glaeser et al., 2001). Empirical research in marketing and urban economics reaches similar conclusions: assortment breadth shapes consumer choice, and product variety rises systematically with market size (Hoch et al., 1999; Schiff, 2014). When local variety is limited, residents rationally spend elsewhere to access novelty, niche, or status goods unavailable in-network.

Secondly, upstream suppliers are often non-local. Even when final sales occur locally, many merchants source inputs (wholesale goods, packaging, software, logistics) from outside the area. In I–O terms, high import coefficients in supply chains channel local demand into purchases from other regions, mechanically shrinking the multiplier (Hughes, 2018; IMPLAN, 2023).

Thirdly, payment design can accelerate exit. If a local instrument is redeemable 1:1 into national money or can be converted without friction, its marginal propensity to redeem (formalized in §1.4) is high: value quickly returns to the outside payment system unless geographic or merchant restrictions bind. Evidence from municipal voucher programs, such as South Korea’s regional gift certificates and Taiwan’s Triple Stimulus Vouchers, shows short-run spending boosts concentrated where acceptance and discounts are strong, but limited persistence once the program window closes. This pattern is consistent with rapid reversion to national currency when constraints relax (Kim and Oh, 2021; Hua, 2022).

Two empirical tools help trace these dynamics. Local Multiplier 3 (LM3) analysis measures how much of an initial euro stays and recirculates through two or three rounds of spending (LM3/NEF, n.d.). Input–output decompositions, such as IMPLAN-style models, identify which channels (imports, profits, or taxes) dominate outflows and where procurement shifts or institutional design could most effectively “plug” the leaks (Hughes, 2018; IMPLAN, 2023).

1.2.2 The “forced-barter” intuition and its limits

A common intuition is that if a token can only circulate locally, the community effectively recreates barter: trade becomes indirect, but value cannot escape, so matching should thicken. Search-theoretic monetary models clarify both the promise and the limits of this idea. In settings with double-coincidence and search frictions, a widely accepted medium of exchange dominates direct swaps by

compressing search costs and coordinating expectations (Kiyotaki and Wright, 1993; Lagos and Wright, 2005). A closed-loop token can therefore outperform barter while constraining exit routes, but only under specific conditions.

Payment systems rely on indirect network effects: consumers value a token that many merchants accept, and merchants accept it because many consumers use it (Rochet and Tirole, 2005; Rysman, 2007). Without sufficient acceptance density on both sides, the option value of waiting remains high, velocity stays low, and the system struggles to reach self-sustaining circulation. Conversion routes also matter. If balances can be redeemed at par or cheaply converted into national currency, the system leaks by design. Evidence from voucher programs shows this trade-off clearly: discounts and ease of use raise adoption, but frictionless redemption raises the propensity to exit unless geographic or merchant restrictions are tight (Kim and Oh, 2021; Hua, 2022).

Even with dense acceptance and limited exit, participants must also believe that future spending opportunities will be plentiful and timely. When inventories are thin or restocks uncertain, agents rationally hold balances in anticipation of better matches - behavior consistent with models where liquidity provides option value. Demurrage or expiry can lower the return to waiting, but their impact depends on both acceptance and supply; case evidence from the Chiemgauer confirms this sensitivity (Gelleri, 2009; Hayashi, 2021).

Put together, “forced barter” through a closed token only works when acceptance density passes a tipping threshold, exit routes are narrow enough that value must recirculate, and participants expect a credible menu of locally demanded goods and services. When any of these conditions weakens, closed tokens tend to accumulate idle balances rather than facilitate active exchange.

1.3 Utility of money per se

Beyond the canonical liquidity story, a broader set of behavioral mechanisms can make tokens privately valuable within small loops and lengthen holding periods even when designers intend fast recirculation. Uncertainty about future spending opportunities creates a strong incentive to delay exchange -participants often hold liquidity as insurance. The precautionary or buffer-stock motive captures this logic: agents build liquid buffers to smooth consumption and remain ready for rare opportunities (Deaton, 1991; Carroll, 1997). When on-network purchases require tokens, the buffer stock is held in tokens rather than euros. A higher perceived volatility of arrivals, such as irregular

restocks, therefore increases token balances and prolongs saving spells, while publishing predictable restock calendars should lower balances by reducing uncertainty (Carroll, 1997). When spending opportunities are uncertain and search requires effort, holding tokens also carries an option value. The more unpredictable the future menu of goods, the more valuable it becomes to wait for a better match (Nosal and Rocheteau, 2011). People act as if patience itself were an asset: they delay spending now because holding preserves the chance of a better opportunity later. In this sense, individuals “lend time” by keeping balances. Providing clear information about supply cadence, and reducing search costs, can reduce uncertainty without eliminating anticipation entirely.

In many token systems, spending does not adjust smoothly but clusters around clear thresholds. Effort and patience increase nonlinearly as people approach a goal. Near a threshold, motivation spikes, and even a small head start can reinforce persistence (Kivetz, Urminsky, and Zheng, 2006; Nunes and Drèze, 2006). In token loops, participants may save until they reach a clear target, such as holding enough credit for a desired item, stretching holding periods and producing lumpy bursts of spending once the target is achieved.

Individuals mentally separate resources by source and purpose (Thaler, 1999). Token balances can be placed in dedicated mental accounts for merchandise, food, or activities, which stabilizes spending patterns but can also foster idle balances when categories remain inactive. The pain of paying also matters. When payment is immediate and salient, willingness to spend drops; when it is delayed or abstracted, as with prepaid balances, spending becomes easier (Prelec and Loewenstein, 1998; Soman, 2001; Prelec and Simester, 2001). Prepaid tokens can therefore increase spending by decoupling payment from consumption, though if acquiring tokens is itself effortful or emotionally costly, users may hoard them instead.

People treat numerical units as meaningful thresholds. Very small denominations can feel inconsequential and be spent casually, while round-number prices can increase psychological barriers (Raghubir and Srivastava, 2009). Conversely, large denominations can create friction by making it harder to “break” balances, thereby increasing idle holdings. Designers can influence spending rhythms by setting price granularity and ensuring that common purchases do not require holding large balances.

In two-sided token systems, participation on one side depends on activity on the other. If suppliers expect low demand, they list fewer attractive goods; if buyers expect thin supply, they

conserve balances. This coordination problem resembles platform adoption dynamics and can generate multiple equilibria (Rochet and Tirole, 2003). Even without explicit fees, thin network density can therefore prolong holding and reduce circulation until confidence in ongoing activity is established.

Hoardability can be defined as the marginal utility of holding an additional token between periods when immediate spending is feasible. In this paper it is defined as factor γ . Intuitively, γ summarizes the combined benefits of liquidity services highlighted in standard monetary models, precautionary motives, real-option value, goal dynamics, and mental accounting. A higher γ indicates a stronger incentive to carry balances forward and, all else equal, lower transaction velocity. This concept is particularly relevant in non-redeemable systems. When conversion into national currency is ruled out, the key decision margin becomes spend now or hold for later. If participants anticipate future restocks, better matches, or more attractive bundles, γ increases; if acceptance is dense and desirable low-priced items are reliably available in the present, γ declines.

A practical implication is that γ can be lowered by raising immediate usability: increasing acceptance points, ensuring predictable micro-restocks, and maintaining a steady supply of low-ticket items all reduce the option value of waiting.

1.4 From micro to macro: how MPR eats MPC

1.4.1 A minimal accounting frame

We begin from the standard expenditure identity at the household/agent level: given one unit of period income (or token receipts) that is available for disposition, the shares allocated to current consumption and saving must sum to one, i.e.,

$$1 = MPC + MPS$$

where MPC is the marginal propensity to consume and MPS is the marginal propensity to save (textbook Keynesian accounting; see, e.g., Blanchard & Johnson, 2013).

When a convertible/redeemable alternative currency is present, additional margin appear: redeeming/convertibility back to base money. Considering period- t token receipts normalized to one. Their uses partition into three mutually exclusive shares:

$$1 = \underbrace{MPC}_{\text{current consumption with the token}} + \underbrace{MPS}_{\text{carry/hoard token to } t+1} + \underbrace{MPR}_{\text{redeem/convert to national money}} .$$

This accounting is not just bookkeeping; it pins down the macro stakes. MPC maps into on-network transactions and velocity; MPS captures intertemporal carry tied to the hoardability motive γ ; MPR is the cash-out channel that removes purchasing power from the circuit. Holding MPS fixed, increasing MPR mechanically decreases MPC. Redeemability therefore is not neutral: it embeds a leakage route that competes directly with trade.

What determines each share? Three micro levers introduced earlier do most of the work: acceptance density (how many counterparties, how connected), redemption frictions (par convertibility, fees, speed), and hoardability γ (the continuation value of holding). Dense acceptance shrinks search frictions and raises spend-now incentives (Easley & Kleinberg, 2010; Jackson, 2010; Newman, 2010). Platform pricing and cross-side participation govern whether systems reach a tipping scale (Armstrong, 2006; Caillaud & Jullien, 2003; Parker & Van Alstyne, 2005). Monetary search and laboratory evidence explain when the value of holding money coordinates trade and when, in thin networks, too much holding depresses matches (Nosal & Rocheteau, 2011; Camera & Casari, 2014; Duffy & Puzzello, 2014).

1.4.2 What redeemability does in practice (evidence by archetype)

For redeemable currencies, par convertibility lowers adoption frictions, users can price 1:1 and “sleep at night”, but it also elevates MPR. The Bristol Pound achieved strong brand visibility but limited durable procurement change; qualitative evaluations emphasize institutional constraints, thin acceptance, and the availability of a sterling exit (Johnson & Harvey-Wilson, 2018; Marshall, 2018). BerkShares (since 2006) show no detectable macro effect on county business dynamics in a synthetic-control design, consistent with soft incentives being outweighed by redemption and by supply chains anchored outside the network (Matti & Zhou, 2022). In the accounting, both cases imply high MPR, keeping MPC low over the horizons that matter for recirculation.

Municipal / national vouchers reliably lift contemporaneous spending, especially with broad acceptance and sizable discounts (Kim & Oh, 2021; Hua, 2022). But the instrument is designed for quick cash-out; as windows close, spending reverts. In our frame: policy temporarily pushes up MPC via subsidies, but structurally high MPR limits recirculation, so multiplier-type effects decay as balances convert back to fiat.

Mutual credit networks, being non-redeemable by design, remove the cash-out door: $MPR \approx 0$. Balances are “redeemed” only via purchases from members. Evidence links such architectures to stabilization and countercyclical use among SMEs and to liquidity-saving through multilateral clearing (Stodder & Lietaer, 2016; Fleischman & Dini, 2020). In our identity, setting $MPR = 0$ expands the feasible share for MPC (subject to MPS) if acceptance is curated and supply chains are represented so that waiting is not privately optimal.

Demurrage seeks to tax MPS (lower γ) and raise MPC, but its success hinges on acceptance and menu credibility. When density is thin or inventories are unreliable, demurrage can backfire - agents may try to escape holding costs by converting (where possible) or disengaging (Gelleri, 2009; Hayashi, 2021). In redeemables, adding demurrage often shifts mass from MPS \rightarrow MPR (leakage). In non-redeemables with thin acceptance, it may simply compress balances without producing additional trades. In conclusion, redeemability guarantees an exit ($\uparrow MPR \Rightarrow \downarrow MPC$). Non-redeemability removes the exit but requires active management of MPS via acceptance density and predictable supply. Demurrage can help only after those are credible.

1.4.3 From levers to outcomes: linking mechanisms to behavior

Each design lever in a local token system ultimately affects how people spend, save, or redeem their balances. Think of it as a chain reaction: small shifts in network structure or user experience flow through to the core behavioral margins, how much is spent (MPC), saved (MPS), or redeemed (MPR).

When acceptance density rises (tokens can be used in more places) participants spend less time searching for an outlet and feel more confident about finding a good match. This lowers search costs and raises the expected quality of transactions (Easley & Kleinberg, 2010; Jackson, 2010; Newman, 2010). As a result, people are more willing to spend (MPC increases) and less inclined to hold onto balances (MPS falls).

Lower redemption frictions, such as instant payouts or zero-fee conversions, have a different effect. They encourage merchants to redeem faster, raising the redemption share (MPR). Occasionally, such changes also boost short-term spending at eligible outlets, especially when coupled with discounts or subsidies, but the value tends to exit the system soon after (Kim & Oh, 2021; Hua, 2022).

High hoardability (γ) caused by few accepting merchants, infrequent restocks, or precautionary motives - works in the opposite direction. Participants treat tokens more like savings, raising MPS and reducing spending. MPR may not change much unless large token balances eventually push holders to cash out (Nosal & Rocheteau, 2011; Duffy & Puzzello, 2014). These dynamics explain familiar patterns: a “J-shaped” spike and fade in voucher schemes, the gradual and lasting build-up seen in successful mutual credit systems, and the fragility of redeemable tokens in thin markets.

1.4.4 Measuring and diagnosing token performance

Most applied studies of complementary or local currencies rely on multi-merchant transaction data or issuer ledgers that allow acceptance, redemption, and circulation to be measured in detail. While such datasets enable sophisticated network analysis, the underlying measurement logic travels well across settings. Across the literature, evaluation repeatedly centers on four core dimensions: acceptance, redemption, holding (or hoarding), and leakage, as the primary channels through which design choices shape circulation outcomes.

Acceptance is typically measured through the size and structure of the acceptance network. Some studies count active acceptance points, while others construct buyer–seller graphs and summarize their connectivity using standard network statistics such as degree, component structure, or path length (Easley and Kleinberg, 2010; Jackson, 2010; Newman, 2010). Where richer network data are available, concentration and centrality measures are used to assess whether exchange relies on a small number of hubs and to identify which nodes would most expand reach if onboarded (Borgatti et al., 2009; Opsahl, Agneessens, and Skvoretz, 2010).

When convertibility is permitted, redemption is commonly summarized either by a marginal propensity to redeem (MPR) or by “time outstanding” measures that track the age distribution of balances. Both are directly observable from issuer records. In settings without complete ledgers, studies sometimes complement administrative data with structured surveys or panel data, typically reported with explicit recall windows to limit measurement error.

Holding or hoarding behavior is usually inferred from balance dynamics. High end-of-period balances relative to inflows suggest that units are being retained rather than cycled. Timing measures, such as the delay between receipt and first expenditure, are also used to characterize how quickly currency circulates under different acceptance conditions.

Finally, leakage, value that exits the intended circuit, is often examined using local multiplier logic, such as LM3-style tracking of successive spending rounds, and input–output reasoning about import shares and external purchases (New Economics Foundation, 2002; Miller and Blair, 2009). In redeemable systems, leakage is also directly observable through conversions back into national currency.

1.4.5 Summary:

The MPC/MPS/MPR identity is the macro microscope for small-network currencies. Redeemability raises MPR and, holding MPS fixed, eats MPC. Non-redeemability removes the exit but leaves MPS to manage; acceptance density and predictable menus lower γ and unlock MPC. Demurrage taxes MPS and works only after density and menus are credible; otherwise it diverts balances into MPR (if redeemable) or into off-network purchases (if not). Designing with these levers and measuring each margin turns scattered case studies into a coherent policy toolkit (Armstrong, 2006; Borgatti et al., 2009; Easley & Kleinberg, 2010; Jackson, 2010; Newman, 2010; Nosal & Rocheteau, 2011; Camera & Casari, 2014; Duffy & Puzzello, 2014; Stodder & Lietaer, 2016; Fleischman & Dini, 2020; Gelleri, 2009; Hayashi, 2021; Kim & Oh, 2021; Hua, 2022; Matti & Zhou, 2022).

2. CONCEPTUAL FRAMEWORK

2.1 Introduction

Consider a town populated by many artisans and many potential trading partners. Exchange centers on paintings and sculptures, goods whose value is highly subjective. Although gains from trade exist, random encounters rarely produce mutually acceptable swaps, because preferences must align closely on both sides.

Begin with barter. A painter may meet a sculptor, yet most meetings fail to produce exchange because both parties must want exactly what the other offers at the same time. With taste-based goods, this coincidence of wants is especially thin. After repeated failures, a second friction becomes salient. Participation is costly. Producing additional work and showing up to trade require time and effort. When the likelihood of a successful swap is low, artisans rationally reduce what they bring to market or stop attending altogether. Attendance falls, the market shrinks, and matching becomes even less likely. The problem is self-reinforcing.

Now introduce a neutral token. The town adopts a standardized medium of exchange that is accepted by convention among participants and cannot be redeemed outside the local system. Sale and purchase no longer need to be synchronized. A painter can sell to any buyer who values the work and holds tokens, then later use those tokens to buy a sculpture from a different seller. Relative to barter, the set of feasible trades expands and the expected return to participation rises. When sales become more likely, producing and bringing goods is worth the effort. Higher attendance increases variety, and greater variety further raises the probability that someone finds a suitable match. Trade volume increases relative to the barter benchmark.

Over time, however, holding tokens becomes privately attractive. When goods are subjective and supply is irregular, tokens provide flexibility. Participants keep balances because they do not know what will appear in future market rounds and prefer to preserve the option to buy when a better match becomes available. Strategic waiting reinforces this behavior. If artisans expect that later sessions might feature more appealing works, they postpone spending and retain tokens. The token continues to function as a medium of exchange in principle, but a meaningful share of purchasing power is withheld from current transactions.

Once hoarding becomes sufficiently widespread, market activity slows. With fewer tokens being spent, effective demand for paintings and sculptures declines. Sellers face a lower probability of sale and respond by bringing fewer works or attending less frequently, since participation is costly. Reduced attendance lowers variety, and lower variety further increases the incentive to wait. The system still has money in a narrow sense, yet exchange in real goods contracts because balances are held rather than used. In this setting, money can resolve the matching problem relative to barter and still coincide with low trade volume when it is valued as a store of value.

Fundamentally token systems succeed or fail along a small number of margins: can people spend them (acceptance), do they want to hold them (γ), and can they cash out (redemption). These ideas are intuitive, but they are easy to misstate when moving between “pure exchange” logic, where total quantities are fixed, and market-demand settings where willingness to transact matters and supply responds to expected sales, providing a short formal bridge by translating the picture into a stripped-down model with two real goods, A and B, and a monetary asset, M. Trading volume is defined as the value of exchange in real goods, and hoarding is captured in reduced form by allowing agents to allocate a positive share of resources to holding M rather than spending all purchasing power on A and B. Under Cobb–Douglas preferences, this delivers a clean “hoarding channel”: when money enters utility with a positive weight, Marshallian demands for non-money goods shrink proportionally, so aggregate expenditure and exchange in A and B fall, even though money relaxes barter-style matching constraints. Proof also clarifies why an Edgeworth-box or contract-curve analysis does not generate this result by itself, and it briefly outlines extensions, including cash-in-advance constraints and non-hoardable (demurrage) currency, under which allocations and trading outcomes shift in the expected direction.

2.1.1 Preliminaries and notation

We are assuming two-good economy, A and B, with prices $p_A > 0$ and $p_B > 0$. Later a third good is introduced (M), denotes a utility-bearing, hoardable money with price $p_M > 0$.

Baseline preferences: Without M, agent i has Cobb–Douglas utility

$$U_i(A_i, B_i) = A_i^{\alpha_i} \cdot B_i^{\beta_i} \text{ with } \alpha_i \in (0,1), \beta_i \in (0,1) \text{ and } \alpha_i + \beta_i = 1$$

Preferences (with money and hoarding): With M present, the utility is $U'_i(A'_i, B'_i, M_i) = A_i'^{\alpha'_i} \cdot B_i'^{\beta'_i} \cdot M_i^{\gamma_i}$, where $\gamma_i \in (0,1)$ is the utility share of money and

$$\alpha'_i + \beta'_i + \gamma_i = 1 \Rightarrow \alpha'_i + \beta'_i = 1 - \gamma_i$$

To isolate the hoarding channel assuming no preference changes between A and B

($\frac{\alpha_i}{\beta_i} = \frac{\alpha'_i}{\beta'_i}$) therefore $\alpha'_i = k \cdot \alpha_i, \beta'_i = k \cdot \beta_i$, for some scaling factor $k > 0$

From total share condition: $k \alpha_i + k \beta_i = k (\alpha_i + \beta_i) = k \cdot 1 = 1 - \gamma_i \Rightarrow k = 1 - \gamma_i$

$\alpha'_i = (1 - \gamma_i)\alpha_i$ and $\beta'_i = (1 - \gamma_i)\beta_i$, so the A:B taste ratio is unchanged and only the total share on real goods falls to $(1 - \gamma_i)$.

Budget constraints: Choices satisfy $p_A A_i + p_B B_i \geq I_i$, if preferences are monotonic (“more is better”), this will actually bind as an equality $p_A A_i + p_B B_i = I_i$ (without money) and $p_A A'_i + p_B B'_i + p_M M_i \leq I_i$ when M is available. All quantities are non-negative.

Endowments and incomes: agent i has an endowment vector $\omega_i = (\omega_{iA}, \omega_{iB}, \omega_{iM}) \in \mathbb{R}_+^3$
Aggregate endowments are $\bar{A} = \sum_i \omega_{iA}, \bar{B} = \sum_i \omega_{iB}, \bar{M} = \sum_i \omega_{iM}$

At prices $p = (p_A, p_B, p_M)$ income of agent i is $I_i = p_A \cdot \omega_{iA} + p_B \cdot \omega_{iB} + p_M \omega_{iM}$
In our empirical setting we say that $\omega_{iM} = 0$, so $I_i = p_A \cdot \omega_{iA} + p_B \cdot \omega_{iB}$

2.1.2 Edgeworth efficiency with a separable money good

Edgeworth box (or three-good prism in this case) cannot, by itself, show that introducing hoardable, utility-bearing money reduces consumption of real goods. Under monotone preferences (“more is better”) and with money separable from A and B, every Pareto-efficient allocation exhausts the aggregate endowments of all goods; adding money therefore leaves the set of efficient A–B allocations unchanged. What does change, once prices and budgets are introduced, is the distribution of A and B across agents: with Cobb–Douglas preferences, giving agent i a money share γ_i scales down that agent’s real-goods expenditure by the non-money share $(1 - \gamma_i)$. Holding real-goods taste parameters (α_i, β_i) and incomes fixed, competitive equilibrium selects a point along the original contract set in which each agent’s share of A and B is proportional to these non-money weights; in the symmetric benchmark, only the relative weights $\gamma_1 : \gamma_2$ matter. Next subsection formalizes these statements and derives the corresponding allocation formulas.

Claim (invariance of the A–B projection): Two agents have monotone Cobb–Douglas preferences over two goods (A, B) and utility-bearing money M:

$U_i(A_i, B_i, M_i) = A_i^{\alpha_i} \cdot B_i^{\beta_i} \cdot M_i^{\gamma_i}$, with $\alpha_i > 0, \beta_i > 0, \gamma_i > 0$, and $\alpha_i + \beta_i + \gamma_i = 1$.

Aggregate resources are $(\bar{A}, \bar{B}, \bar{M})$, with feasibility:

$$A_1 + A_2 = \bar{A} \quad B_1 + B_2 = \bar{B} \quad M_1 + M_2 = \bar{M}$$

A weighted-utilitarian planner maximizes:

$$\lambda \cdot (\alpha_1 \ln A_1 + \beta_1 \ln B_1 + \gamma_1 \ln M_1) \\ (1 - \lambda) \cdot (\alpha_2 \ln A_2 + \beta_2 \ln B_2 + \gamma_2 \ln M_2)$$

subject to the three feasibility constraints.

Result (first-order conditions):

$$\frac{A_1}{A_2} = \frac{\lambda \cdot \alpha_1}{(1 - \lambda) \cdot \alpha_2} \quad \frac{B_1}{B_2} = \frac{\lambda \cdot \beta_1}{(1 - \lambda) \cdot \beta_2} \quad \frac{M_1}{M_2} = \frac{\lambda \cdot \gamma_1}{(1 - \lambda) \cdot \gamma_2}$$

Conclusion. The efficient split of A depends only on $(\alpha_1, \alpha_2, \lambda)$; the efficient split of B depends only on $(\beta_1, \beta_2, \lambda)$. The money exponents (γ_1, γ_2) only affect the split of M. With monotone preferences, all efficient allocations exhaust $(\bar{A}, \bar{B}, \text{and } \bar{M})$. Therefore, adding a separable money good does **not** change the set of efficient A–B allocations (the projection of the Pareto set onto the A–B plane).

2.1.3 Distributional effects in competitive equilibrium (budgets)

Normalize $p_M = 1$ (money is the numeraire). Prices of real goods are (p_A, p_B) . For agent i , income equals the market value of the endowment:

$$I_i = p_A \cdot \omega_{iA} + p_B \cdot \omega_{iB} + \omega_{iM}$$

If we give money a positive utility share γ_i and keep each agent's A:B taste ratio unchanged by shrinking real-goods shares proportionally: effective shares on A and B become $(1 - \gamma_i) \cdot \alpha_i$ and $(1 - \gamma_i) \cdot \beta_i$

Marshallian demands:

$$A'_i = \frac{(1-\gamma_i) \cdot \alpha_i \cdot I_i}{p_A} \quad B'_i = \frac{(1-\gamma_i) \cdot \beta_i \cdot I_i}{p_B} \quad M'_i = \gamma_i \cdot I_i$$

2.1.4 Symmetric benchmark (closed form)

Assumptions:

1. Identical real-goods taste proportions before money is introduced:

$$\alpha_1 = \alpha_2 = \alpha, \beta_1 = \beta_2 = \beta \text{ with } \alpha + \beta = 1$$

2. Equal incomes: $I_1 = I_2$

3. Market clearing: $A'_1 + A'_2 = \bar{A}$ and $B'_1 + B'_2 = \bar{B}$.

Proposition: Relative preference weights on real goods equivalently, the non-money share ($1 - \gamma$) determine the division on A and B

$$\frac{A'_1}{A'_2} = \frac{(1 - \gamma_1)}{(1 - \gamma_2)} \quad \frac{B'_1}{B'_2} = \frac{(1 - \gamma_1)}{(1 - \gamma_2)}$$

Hence:

$$A'_1 = \frac{1 - \gamma_1}{1 - \gamma_1 + 1 - \gamma_2} \cdot \bar{A}$$

$$B'_1 = \frac{1 - \gamma_1}{1 - \gamma_1 + 1 - \gamma_2} \cdot \bar{B}$$

(And $A'_2 = \bar{A} - A'_1$, $B'_2 = \bar{B} - B'_1$)

Proof :

$$A'_1 = \frac{(1 - \gamma_1) \cdot \alpha \cdot I}{p_A} \text{ and } I_1 = I_2 = I$$

$$\frac{A'_1}{A'_2} = \frac{\frac{(1 - \gamma_1) \cdot \alpha \cdot I}{p_A}}{\frac{(1 - \gamma_2) \cdot \alpha \cdot I}{p_A}} = \frac{1 - \gamma_1}{1 - \gamma_2}$$

We could impose $A'_1 + A'_2 = \bar{A}$ to recover exact values. B is analogous.

Comparative statics:

if $\gamma_1 > \gamma_2$, then $(1 - \gamma_1) < (1 - \gamma_2)$, so $A'_1 < A'_2$ and $B'_1 < B'_2$. Logically the agent with the larger money preference ends up with a smaller fraction of both real goods. In this symmetric benchmark, only the ratio $(1 - \gamma_1):(1 - \gamma_2)$ matters, as α , β , p_A and p_B cancels.

2.1.5 General ceteris paribus (unequal incomes and/or real-goods tastes)

Without symmetry, the division is governed by income-weighted, taste-weighted, non-money shares. Assuming we have two agents, before we introduce money:

$$\text{Agent } i \in \{1,2\}, u_i = A^{\alpha_i} \cdot B^{\beta_i} \text{ with } \alpha_i + \beta_i = 1$$

Planner maximizes:

$$\max_{(A_1, A_2, B_1, B_2)} \lambda \ln u_1 + (1 - \lambda) \ln u_2$$

subject to:

$$A_1 + A_2 = \bar{A}, B_1 + B_2 = \bar{B}.$$

Lagrangian:

$$L = \lambda, [a_1 \ln A_1 + b_1 \ln B_1] + (1 - \lambda), [a_2 \ln A_2 + b_2 \ln B_2] + \mu_A(\bar{A} - A_1 - A_2) + \mu_B(\bar{B} - B_1 - B_2)$$

FOC for A_1, A_2

$$\begin{aligned} \frac{\partial L}{\partial A_1}: \quad \frac{\lambda \alpha_1}{A_1} - \mu_A &= 0 \quad \Rightarrow \mu_A = \frac{\lambda \alpha_1}{A_1} \\ \frac{\partial L}{\partial A_2}: \quad \frac{(1 - \lambda) \alpha_2}{A_2} - \mu_A &= 0 \quad \Rightarrow \mu_A = \frac{(1 - \lambda) \alpha_2}{A_2} \end{aligned}$$

Equalize μ_A from both FOCs.

$$\frac{\lambda \alpha_1}{A_1} = \frac{(1 - \lambda) \alpha_2}{A_2} \quad \Rightarrow \frac{A_1}{A_2} = \frac{\lambda \alpha_1}{(1 - \lambda) \alpha_2}$$

Same logic for B_1, B_2 :

$$\frac{\lambda\beta_1}{B_1} = \frac{(1-\lambda)\beta_2}{B_2} \Rightarrow \frac{B_1}{B_2} = \frac{\lambda\beta_1}{(1-\lambda)\beta_2}$$

After adding money:

$$u_i = A^{\alpha'_i} \cdot B^{\beta'_i} \cdot M^{\gamma_i} \text{ with } \alpha'_i + \beta'_i + \gamma_i = 1$$

We assume that ratio of preferences towards goods A and B didn't change after money was introduced (ceteris paribus)

$$\frac{\alpha_i}{\beta_i} = \frac{\alpha'_i}{\beta'_i} = k_i$$

From the ratio and sum to one condition we get:

$$\alpha'_i = (1-\gamma) \cdot \alpha_1 \text{ and } \beta'_i = (1-\gamma_i) \cdot \beta_i$$

Use the standard social planner with weight $\lambda \in (0,1)$ and two agents we get:

$$\max \lambda \ln \cdot u_1 + (1-\lambda) \ln \cdot u_2$$

subject to resource constraints:

$$A_1 + A_2 = \bar{A} \quad B_1 + B_2 = \bar{B}, \quad M_1 + M_2 = \bar{M}$$

With Cobb-Douglas, $\ln u_i = \alpha'_i \ln A'_i + \beta'_i \ln B'_i + \gamma_i \ln M_i$

FOCs give good-to-good split ratios:

$$\frac{A'_1}{A'_2} = \frac{\lambda\alpha'_1}{(1-\lambda)\alpha'_2}, \quad \frac{B'_1}{B'_2} = \frac{\lambda\beta'_1}{(1-\lambda)\beta'_2}, \quad \frac{M_1}{M_2} = \frac{\lambda\gamma_1}{(1-\lambda)\gamma'_2}$$

By plugging in

$$\alpha'_i = (1-\gamma) \cdot \alpha_1 \text{ and } \beta'_i = (1-\gamma_i) \cdot \beta_i$$

we get:

$$\frac{A'_1}{A'_2} = \frac{\lambda\alpha_1}{(1-\lambda)\alpha_2} \cdot \frac{1-\gamma_1}{1-\gamma_2}$$

Hence holding the A/B taste ratios fixed for each agent, introducing M doesn't change the A–B tradeoff shape; it uniformly tilts who gets A and B by the single scalar $\frac{1-\gamma_1}{1-\gamma_2}$, basically same answer as with symmetric preferences.

Conclusion: In the Edgeworth box framework, we are in a pure exchange economy with a fixed aggregate endowment of each good. The total quantities of are given exogenously, and every Pareto-efficient allocation must exhaust these totals. “More is better” preferences ensure that no goods are left unconsumed: if one agent voluntarily reduces their consumption of a real good in favour of holding money, another agent will necessarily absorb the surplus. In this setting, the introduction of a utility-bearing money good can only affect the distribution of real goods between agents, never their aggregate quantities. This is a structural feature of the model - the Edgeworth box literally has no mechanism for goods to disappear or remain unsold.

Edgeworth would only predict less volumes of good A and B for both agents only if money entered feasibility (CIA/reserve) or used resources; our design does neither, hence we analyze the demand channel and show why non-hoardability removes the leak.

2.2 A Market-Demand Proof: Utility-Bearing Money Reduces Real-Goods Demand

In contrast, the market-demand approach operates in a setting where goods are produced in response to demand, and quantities are not fixed *a priori*. Prices and incomes are taken as given, and producers supply only what buyers wish to purchase at those prices. In this framework, if agents allocate part of their expenditure to a hoardable money good, the total demand for A and B falls. Since output is demand-driven, lower aggregate demand translates into lower production and, therefore, lower total real-goods consumption. The “missing” goods from one agent's budget are not automatically reassigned to someone else; they simply are not produced in the first place. As a result, the introduction of a utility-bearing money good leads to a contraction in real-goods market activity, even at unchanged prices.

We work at arbitrary prices (p_A, p_B, p_M) and income:

$$I_i = p_A \cdot \omega_{iA} + p_B \cdot \omega_{iB}$$

Proposition: Hoarding steals real-goods demand

Under the re-parameterized Cobb–Douglas preferences

$$U'_{i(A'_i, B'_i, M_i)} = A_i^{(1-\gamma_i)\alpha_i} \cdot B_i^{(1-\gamma_i)(1-\alpha_i)} \cdot M_i^{\gamma_i}$$

Marshallian demands satisfy, for any

(p_A, p_B, p_M) and $I_i > 0$:

$$A'_i = \frac{(1-\gamma_i) \cdot \alpha_i \cdot I_i}{p_A}, \quad B'_i = \frac{(1-\gamma_i) \cdot (1-\alpha_i) \cdot I_i}{p_B}, \quad M_i = \frac{(\gamma_i I_i)}{p_M}$$

Proof:

For Cobb–Douglas utility with exponents summing to one, Marshallian demands allocate fixed expenditure shares by the exponents. Because $\alpha'_i = (1 - \gamma_i) \cdot \alpha_i$ and $\beta'_i = (1 - \gamma_i)(1 - \alpha_i)$, the expenditure share on A is α'_i , on B is β'_i , and on M is γ_i . Hence

$$A'_i = \frac{(1-\gamma_i) \cdot \alpha_i \cdot I_i}{p_A}, \quad ; \text{ similarly for } B'_i; \text{ and } M_i = \frac{(\gamma_i I_i)}{p_M}.$$

Corollary: Lower trading volume in A and B at given prices

The value of trade in the real goods (aggregate expenditure on A and B) equals

$$V'(p) \equiv \sum_i (p_A A'_i + p_B B'_i) = \sum_i (1 - \gamma_i) I_i .$$

If $\gamma_i \in (0,1)$ for at least one agent, then $V'(p) < \sum_i I_i = V(p)$, the baseline without M.

Equivalently, the aggregate demand for A and B shifts inward pointwise at any price.

Interpretation: Giving money a utility share γ_i diverts a γ_i fraction of each agent's budget away from A and B. With Cobb–Douglas, that simply multiplies each agent's real-goods demands by $(1 - \gamma_i)$ independent of prices. Hence aggregate real-goods demand shifts inward: at fixed prices, quantities fall; if quantities are fixed (vertical supply/endowment), prices fall; with standard upward-sloping supply, both equilibrium quantities (down) and prices (down) typically adjust.

2.3 General-Equilibrium Clarifications

In a pure exchange economy with endowments only in A and B (no initial M), each agent's income is $I_i = p_A \cdot \omega_{iA} + p_B \cdot \omega_{iB}$. At any price vector p , the Marshallian demands derived above hold and aggregate expenditure on A and B equals $(1 - \bar{\gamma}) \cdot \sum_i I_i$, where $\bar{\gamma}$ is the income-weighted average of γ_i . Relative prices will adjust between equilibria with and without M, but the share of

income devoted to A and B is mechanically smaller once $\bar{\gamma} > 0$. Thus any equilibrium comparison at a common price vector shows a strictly lower value of trade in A and B when M carries utility.

Definition (Trading volume).

For empirical or simulation purposes, we define trading volume in real goods at prices p as $V(p) \equiv \sum_i (p_A A_i + p_B B_i)$. Under the hoarding specification, $V'(p) = \sum_i (1 - \gamma_i) I_i$. With no initial endowment of M, $\sum_i I_i = p_A \cdot \bar{A} + p_B \cdot \bar{B}$,

so $V'(p) = (1 - \bar{\gamma}) \cdot (p_A \cdot \bar{A} + p_B \cdot \bar{B})$, which is strictly below the baseline $V(p)$ when any $\gamma_i > 0$.

2.4 Non-Hoardable Currency (Demurrage) and the Reversal

A non-hoardable or demurrage-bearing currency can be modeled by imposing a holding cost $\tau > 0$ per unit of M (or a decay in utility via an effective share $\gamma_i(\tau)$ that tends to zero as τ increases). In the limit as the holding cost eliminates the private value of hoarding ($\gamma_i \rightarrow 0$ in the optimum), the Marshallian demand shares for A and B return to their baseline values and the trading volume in A and B re-expands to the no-M benchmark. This formalizes the intuition that well-designed, non-hoardable currencies can sustain circulation without siphoning purchasing power away from real goods.

2.5 Conclusion

In a standard market (budget-constraint) frame with Cobb–Douglas preferences, adding a utility-bearing, hoardable money good with positive share γ_i rigorously reduces demand for non-money goods in direct proportion to $(1 - \gamma_i)$. Aggregate trading volume in A and B consequently falls at any given price vector, and the economy's demand schedules for real goods shift inward. Edgeworth-box reasoning on its own does not generate this result unless a liquidity or technological link couples M to A and B. These observations justify using a controlled, non-hoardable school currency in the empirical part of the thesis to avoid the hoarding drain on real exchange

3. RESEARCH DESIGN AND EMPIRICAL STRATEGY

3.1 Research design

This study adopts an embedded case study approach to explore the potential and practical challenges of implementing a local currency system in a school setting. The aim was to test whether the introduction of alternative currency could stimulate trade and resource circulation among participants who otherwise may have little incentive to engage in exchange, also to observe how students interact with such a currency in a real-world microenvironment.

The research was inspired by prior experiments such as the Bristol Pound and theoretical foundations related to complementary currencies, particularly those that argue that the presence of a secondary medium of exchange can nudge inactive goods into active circulation. The guiding hypothesis was straightforward: certain items, such as used toys or books, may not hold enough perceived value to be worth selling for euros, but could become tradable within a closed local currency system. In this context, the alternative currency acts as a nudge, encouraging exchanges that otherwise wouldn't occur. The principal research question examined whether a local, non-convertible currency could generate higher levels of trade and participation in a small, closed economic environment.

The case study was selected as the most appropriate methodological framework due to the project's exploratory nature, relatively small scale, and the emphasis on real-time behavioral observation. Unlike surveys or hypothetical scenarios, a case study allowed us to observe emergent patterns and adapt the environment dynamically in response to participant behavior. Moreover, the study functioned as a pilot project, with the long-term vision of scaling or adapting the design into a more formal future research agenda.

The research design leans toward a quantitative orientation. While qualitative components such as observation and verbal interactions played a supportive role, the primary focus was on numerical data: participation rates, transaction frequency, item types, and observable changes following specific interventions.

As the experiment progressed, we introduced intervention in response to observed low circulation. To counteract the decline in currency circulation, a series of collectible cards featuring school-related inside jokes was introduced. These cards, while low in production cost, had perceived

non-monetary value among the students and were intended to increase engagement and encourage spending.

Although the setting was artificial in the sense of introducing a new currency in a school, the behaviors observed were largely natural. Students were not pressured to act in any particular way beyond the structure of the shop and the currency itself. Therefore, the study occupies a space between naturalistic observation and controlled behavioral intervention.

Finally, ethical considerations were addressed through passive informed consent. All students and their parents were informed of the experiment in advance, and participation was entirely optional. No identifying personal data was collected at any stage.

3.2 Setting and participants:

The experiment took place in a private school in Vilnius, Lithuania, which is attended by students from kindergarten to grade 11. The school is split into two physically separate buildings: one for younger students (kindergarten to grade 4) and one for older students (grades 5–11). The experiment was conducted only in the building with older students.

At the time of the study, the school had no twelfth-grade students. Grades 5 to 9 had approximately 20 students each, while grade 10 had around 15 students and grade 11 about 10, resulting in a total pool of roughly 115 students having access to the project.

The shop was set up in a hallway during school breaks and operated for eight weeks (excluding one week of holidays). The researcher's role throughout the experiment was semi-participatory. While not formally acting as a central bank (as the currency supply was not actively managed), we designed the currency system, printed and distributed the currency, operated the shop at intervals, and engaged with students in informal discussions to maintain interest and gather brief data. This blend of passive observation and active presence reflected the dual aim of minimizing interference while ensuring operational continuity. School staff assisted only with basic logistics, such as providing a table. Some teachers supported the project by encouraging students at the beginning, but there was no formal school involvement beyond that.

The implementation allowed for observable variation in how participants engaged with the currency. Participation ranged from immediate exchange behavior to the accumulation of balances. Anticipating potential attempts at counterfeiting, a custom stamp was introduced as a preventive

measure to safeguard the integrity of the currency. These behaviors, while minor, were relevant for understanding how participants interacted with the system's rules and constraints.

3.3 Data collection methods:

Data for this project were collected through a shared spreadsheet managed primarily by the student volunteer who operated the school shop. The spreadsheet recorded basic information for each transaction, including the item name, the buyer's self-reported age group (categorized as "younger" or "older" according to grade level), gender, date of purchase, and, when available, corresponding sales information (seller and time of item submission).

While no personal names were recorded to maintain participant anonymity, the dataset still allows for analysis of general behavioral patterns such as participation by age group, gender differences, and changes in transaction frequency over time, as well as item prices, making it possible to evaluate shop's balance.

In total, the dataset includes a modest number of entries, reflecting limited student engagement. The sparse nature of the data is a limitation; however, it is representative of the real behavioral response to the local currency initiative under the given conditions. No formal surveys or structured interviews were conducted. However, informal verbal feedback from students was noted during the operation of the shop, particularly around which items they found desirable or uninteresting.

3.4 Intervention Design

As this was a pilot experiment, no interventions were planned in advance. The research approach intentionally left room for responsive adjustments based on observed behavior, allowing the design to evolve in real time as problems emerged. This adaptive method was chosen due to the exploratory nature of the study and the lack of prior benchmarks in similar educational settings.

The primary issue that caused intervention was a visible decline in student participation after the initial phase of excitement. During first two weeks, several students donated used items, primarily books and a small number of toys, and a few transactions took place. However, toy inventory quickly sold out, and the shop was left with a disproportionate supply of books, which students showed little interest in buying. This was partly due to the socio-economic context of the school: many students were from well-off households and lived in neighborhoods where book exchange stations ("take a

book, leave a book" kiosks) were common. As a result, second-hand books did not carry enough perceived value to drive continued engagement.

To address this stagnation and the resulting hoarding of the toy currency, several interventions were implemented, varying in scale.

3.4.1 Collectible Cards

The first and primary intervention involved the introduction of collectible cards, which were sold for one unit of the toy currency each. These cards featured inside jokes and humorous references relevant to the school's student culture (Appendix 3). Their purpose was to create subjective value without introducing real material costs or distorting the experiment through external financial input. The production costs were minimal, and the cards could not be used outside the in-game context, aligning with the principle of non-convertibility that underpins most local currency models.

3.4.2 Active Buying Phase (Non-Structural)

A less formal intervention was made directly approaching students who had previously expressed interest in contributing items but failed to do so. These one-on-one conversations served as personal reminders or nudges to follow through. While this approach helped sustain short-term interest and engagement with few students, it did not result in a significant increase in item donations and did not affect the system-wide problem of declining participation.

3.4.3 Ethical Boundaries

While we acknowledged that mechanisms like gambling or lottery-based incentives could effectively stimulate circulation, such tools were ruled out due to ethical considerations in a school setting. These options were discussed theoretically as mechanisms for value creation through anticipation or chance but were not implemented due to their unsuitability for minors.

In summary, the intervention design was intentionally flexible and reactive. The measures taken addressed certain economic imbalances (e.g., hoarding) but failed to solve deeper behavioral inertia, such as the lack of ongoing item contributions. This outcome provides valuable insights into the limits of non-material incentives and the challenges of sustaining engagement in closed local economies without continuous external inputs.

3.5 Data Analysis Strategy

Given the exploratory nature and small scale of the dataset, the empirical analysis is primarily descriptive. The aim is not to make generalizable causal claims, but to document behavioral and accounting patterns that can be compared to the mechanisms discussed in the literature and conceptual framework. The analysis follows the standard diagnostic logic used in studies of parallel currencies: opportunity to spend (acceptance), incentives to wait (holding/hoarding). Loss of circulation (leakage) is not operative by design in this setting because tokens are non-redeemable. Analysis is organized around market outcomes for real goods, currency flows and low-velocity/hoarding proxies, and heterogeneity across participant groups.

3.5.1 Dataset structure and preprocessing

The unit of observation is the item that entered the shop during the study period. After removing one aggregate summary line, the analytical sample consists of 64 distinct entries. For each entry we use: item category, age group and gender of the supplier, purchase price paid by the shop (where applicable), time-to-sale (days between listing and sale, or “not sold”), sale price, and item-level profit/loss. “Collectible cards” are treated as a special category: because cards were produced in effectively unlimited supply and appear in the dataset only when purchased, card rows are interpreted as realized spending events, not as inventory items with a meaningful sale probability.

3.5.2 Descriptive market outcomes

Research summarizes the structure and performance of the shop inventory using simple descriptive statistics and cross-tabulations. Specifically, report counts and category shares describe the purchase-price distribution for non-card goods using standard location and range measures (means, medians, and range), and document sale rates both overall and by category (sold versus unsold). For items that sell, we also summarize time-to-sale (mean/median/range) and compare these timing patterns across categories where doing so is meaningful given the small cell sizes. The results are presented using compact tables and basic figures, including category distributions shown both with and without cards.

3.5.3 Currency flows and low-velocity proxies

To connect shop activity to circulation, we summarize token flows using a simple balance-sheet view. Tokens are issued when the shop buys items from students and recovered when students buy items from the shop. Recovered currency is reported separately for purchases of real goods and

collectible cards, since cards are recorded only when purchased and do not behave like inventory. From these totals, we report the implied net unrecovered outflow at the end of the observation window and the share of issued currency effectively tied up in unsold inventory, valued at acquisition cost. These accounting summaries serve as a practical representation for low velocity/hoarding without individual-level data.

3.5.4 Pre/post intervention comparison (two-phase accounting)

The analysis treats the study as two phases: a pre-intervention period with only real goods, and a post-intervention period after collectible cards were introduced. Phase contrasts are reported descriptively: transaction counts and spending totals by phase, the contribution of card purchases to total transactions and recovered currency, and whether changes reflect higher spending by existing participants or broader participation. Because the intervention was introduced reactively and without randomization, these comparisons are interpreted as patterns rather than causal effects.

3.5.5 Qualitative observations

Informal observations from shop operation and student conversations are used to interpret anomalies (e.g., why some categories stalled, perceived desirability, and reasons for holding balances). These notes support interpretation but are not treated as primary evidence.

3.6 Limitations

This study, while valuable as an exploratory case, faced several limitations that should be acknowledged both in terms of design and implementation. The most significant limitation was a rapid decline in student engagement. Although the total potential sample consisted of more than 100 students, active participation diminished after the initial phase. This likely stemmed from the mismatch between the experimental design and the school's socio-economic context. Students at the selected private school generally lacked motivation to trade low-value items, probably due to many of them had little need for second-hand goods in the first place. The choice of location was primarily driven by access constraints. Implementation outside the host institution proved infeasible, as other schools were generally unwilling to participate without clearly defined benefits for students or staff.

From the data integrity perspective, the recorded transaction log is complete conditional on recorded transactions/participants; however, pricing of items lacked a consistent or objective system. In the absence of an agreed valuation framework, item prices were set intuitively, guided loosely by

comparisons to online second-hand marketplaces such as *Vinted*. While this did not introduce significant distortions, it limited the potential for deeper quantitative analysis of pricing behavior.

In terms of generalizability, the findings are most applicable to private schools in urban Lithuanian settings with comparable student demographics. In educational environments where students face tighter personal budget constraints, the perceived value of in-game currency may be higher, which could alter engagement and trading behavior relative to what was observed here.

Ethical and legal constraints, particularly those related to GDPR, also limited the experiment's depth. No personal data were collected, and participants remained anonymous. While this ensured compliance, it prevented the tracking of individual behavior over time, which would have allowed a richer analysis of engagement, currency accumulation, and transaction patterns. Collecting more detailed data would have required formal parental consent and institutional approval which would be logistically difficult in this setting, where privacy-related objections have historically been strong.

The experimental design itself had several internal constraints. The lack of structured rules around item exchange. For example, allowing students to donate books without any reciprocal demand (led to inventory imbalances and discouraged repeat engagement). In retrospect, a more structured incentive mechanism might have prevented the shop from being flooded with unsellable items.

Finally, the project was constrained by limited human and logistical resources. With only one active volunteer and no external funding, the ability to expand or replicate the model in a larger setting was restricted. Future iterations would benefit from a larger support team and partnerships with schools willing to participate more formally.

Despite these limitations, the experiment provided a meaningful proof of concept and highlighted several key challenges and opportunities in designing closed-currency ecosystems in educational settings.

4. DESCRIPTIVE FINDINGS FROM THE PILOT

4.1 Overview of the experiment and dataset

The field experiment was implemented in a Lithuanian private school as a small, closed marketplace. Students could bring unused household items to a school “shop”. The shop bought these items using a school-specific currency that had no value outside the experiment: tokens could not be exchanged for euros and were not accepted anywhere else, including school. Students could then spend their tokens to purchase items from the shop’s inventory. The experiment unfolded in two phases. In the pre-intervention phase, only “real” goods were traded. After several weeks of low trade volume and visible money hoarding, a post-intervention phase began in which the shop introduced original collectible cards.

The resulting dataset is an item-level panel of all goods that entered the shop during the observation period. After removing one aggregate summary row, the analytical sample consists of 64 distinct items. For each item, the dataset records: category (Book, Toy, Other, Our Cards); owner characteristics: namely age group (Younger vs Older students) and gender (F/M); the purchase price paid by the shop (in units of the school currency), which is missing for cards because they were created at zero cost; the time to sale (in days) and coded as the number of days between listing and sale or as “not sold” if the item remained unsold; the selling price received by the shop; and the item-level profit, (for unsold items, this is simply the negative of the purchase price). Out of the 64 items, 42 were sold, while 22 remained unsold by the end of the experiment. The sample includes 20 books, 13 toys, 6 other items, and 25 instances of collectible cards. In total, the shop paid out 123 units of currency to students for their goods and received 77 units back through purchases, with cards accounting for a sizeable share of these incoming payments.

4.2 Market composition and participation

Participation in the school marketplace was highly uneven across age groups and genders. Out of the 64 recorded items, 55 (86%) were supplied by younger students and only 9 (14%) by older students. Gender-wise, 50 items (78%) came from boys and 14 (22%) from girls. Combining both dimensions, the experiment effectively captured the trading behavior of a relatively narrow but enthusiastic subgroup, primarily younger male students, while older students and girls were present but less active on the supply side.

The inventory itself consisted of a mix of “real” goods and school-branded collectibles. Among the 64 items, books were the largest single category of real goods (20 items, 31% of all entries), followed by toys (13 items, 20%), other items such as small accessories or miscellaneous objects (6 items, 9%), In addition, the dataset contains 25 instances of school collectible cards (“Our Cards”), which account for 39% of all item entries. As noted earlier, these card entries represent actual purchases from an effectively unlimited stock rather than a fixed physical inventory, so they should be interpreted as realized spending opportunities rather than as a finite number of goods available

Figure 4.1.

Composition of completed transactions by category (student-contributed goods only)

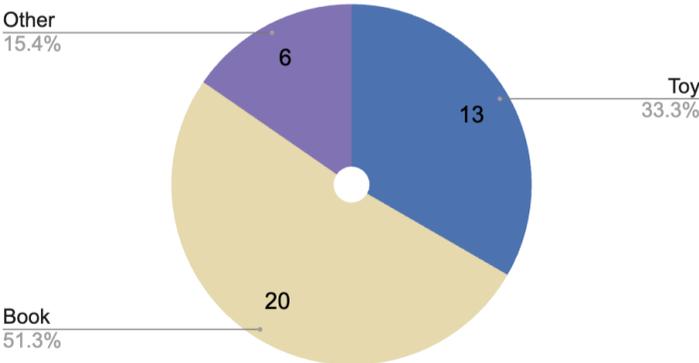
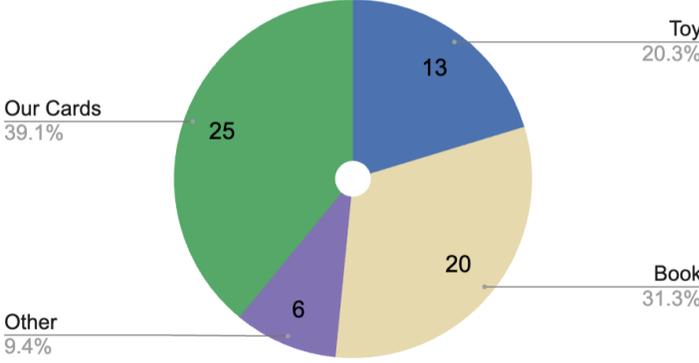


Figure 4.2.

Composition of completed transactions by category (including our cards)



Purchase prices are observed only for non-card goods, since the shop minted cards at zero cost. For the 39 real items (books, toys, and other goods), the purchase price paid to students ranged

from 1 to 8 units of the school currency. The mean purchase price was just over 3 units (3.15), with a median of 3 units; half of all non-card items were bought from students for 2–3 units. Within this, books were typically mid-priced (mean 3.2 units, range 2–6), toys tended to be slightly cheaper on average (mean 2.7 units, range 1–7), and other items somewhat more dispersed (mean 3.8 units, range 1–8). Overall, the shop offered relatively low and tightly clustered purchase prices, reflecting both the modest quality of goods contributed and the intended small scale of the school currency system.

Table 4.1.

Descriptive statistics of listed items by category

Category	Items received	Items sold	Sell rate (%)	Price range	Mean price
All non-card items	39	17	44	1-8	3.15
Books	20	3	15	2-6	3.2
Toys	13	12	92	1-7	2.7
Other items	6	2	33	1-8	4.0

4.3 Trade volume, sales rates, and time to sale

Overall trading activity in the school shop was modest. Focusing on student-contributed goods, a total of 39 distinct items were listed during the experiment, of which 17 were sold, yielding a sale rate of approximately 44%, while 22 items remained unsold by the end of the observation period. Each sold item corresponds to a single transaction, resulting in 17 completed trades in non-card goods. Among books, only 3 of 20 items (15%) were sold, making books the least liquid category in the shop. Toys performed much better: 12 of 13 toys (92%) eventually found a buyer, and 2 of 6 “other” items (33%) were sold. If we focus only on non-card goods, books vs. all other real goods, the contrast is stark: books: 3 sold, 17 unsold - 15% sale rate. Non-book real goods (toys, other items): 14 sold, 5 unsold - 74% sale rate

This pattern suggests that the composition of inventory mattered: playful or more understandably attractive items (toys and some “other” goods) were far more likely to be sold than books. Collectible cards behave differently conceptually, because they were supplied in unlimited quantity and only appear in the dataset when purchased. For cards, therefore, recorded entries should

be interpreted as realized purchases rather than as a finite stock of items with a defined probability of sale.

The timing of trades also reveals a thin and sluggish market. Among the 42 sold items, the time to sale ranged from 0 days (items sold same day of listing) to 52 days, with a mean of about 8.9 days and a median of 8 days. Even items that did eventually sell typically waited around a week or more before finding a buyer. There is some variation across categories: the few sold books changed hands within roughly a week (median 7 days), toys had a similar median time to sale (around 9–10 days) but with much higher dispersion, while the two sold “other” items were either sold immediately or only after several weeks. Taken together, these patterns indicate that trading activity was sporadic, with many items remaining unsold and even successful trades occurring with noticeable delays, consistent with a thin market in which both participation and inventory attractiveness were limited.

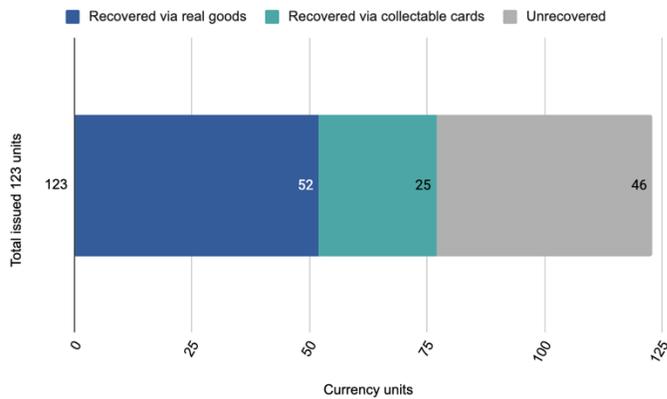
4.4 Currency flows and low velocity in the real-goods segment

From the shop’s perspective, the school currency system can be viewed in two stages: a pre-intervention phase, where only real goods were traded, and a post-intervention phase. Splitting the flows this way helps to separate the underlying liquidity problem from the effect of the intervention. In the first frame, we focus only on the 39 non-card items. To acquire these goods from students, the shop paid out 123 units of the school currency in total. The sold items had a combined acquisition cost of 53 units, while the unsold items cost 70 units.

Reselling the 17 sold items generated 52 units of revenue for the shop. Thus, considering only real goods, the shop recovered 52 out of 123 units it had issued, i.e. about 42% of the injected currency (Figure 4.3). The remaining 71 units represent the sum of the cost of unsold inventory. In aggregate, the real-goods segment of the market produced a net outflow of 71 units from the shop’s balance sheet. For students, this is the mirror image of a substantial net injection of spending power: they were rewarded in tokens for bringing in goods but did not, on average, channel those tokens back into purchases of the available inventory.

Figure 4.3.

Decomposition of issued school currency



Viewed in isolation, the pre-intervention marketplace therefore exhibits a classic thin-market outcome: a configuration where students were generally willing to accept tokens for their contributions (often low-valued or unwanted books) but were much less inclined to spend those tokens on what remained in the shop. Non-book items tended to sell quickly, so the residual inventory was dominated by books that many students did not find attractive. In this environment, tokens were relatively more appealing than the goods still on offer, and both items and currency often remained idle rather than circulating as a medium of exchange.

4.5 Intervention results: collectible cards as a spending nudge

The intervention generated substantial additional trading activity. In total, the dataset records 25 card purchases, each at a price of 1 unit, so cards account for 25 of the 42 transactions in the experiment (around 60%) and 25 of the 77 units spent in the shop (about one third of all observed spending). By construction, the shop’s cost of producing cards was zero, so these 25 units are a pure recapture of previously injected currency. Without the card intervention, the shop would have recovered only 52 units out of 123; with cards, it recovered 77. In other words, cards reduced the net token outflow from 71 to 46 units, partially offsetting the hoarding and trapped-liquidity seen in the non-card segment of the market.

The timing of card purchases also suggests sustained demand once the option existed. For the 25 card entries, the recorded time to sale ranges from 2 to 15 days, with a median of 8 days, similar to the median time-to-sale for other goods but with every recorded card ending in a transaction. Again, this does not mean cards had a 100% “sale rate” in the same sense as finite inventory items, since

their stock was unlimited; rather, it indicates that whenever new card opportunities were presented, students were consistently willing to spend tokens on them within a relatively short period.

Finally, the intervention's distributional pattern is highly concentrated. All 25 card purchases in the dataset are attributed to younger students, and 24 out of 25 to boys. This aligns with the participation patterns on the supply side: younger boys were already the main contributors of items and thus the main initial recipients of tokens. Therefore the cards functioned as a spending nudge for those who already held currency, increasing the velocity of their balances and generating extra transactions, but did not significantly broaden participation to older students or to groups that had not previously acquired tokens. In summary, the collectible cards were effective at turning hoarded balances into observed spending within the active subgroup of participants, while leaving the underlying thinness and inequality of the market largely intact.

4.6 Heterogeneity: age, gender, and item type

In this subsection the focus is on real goods only (books, toys, other items). Collectible cards are excluded from the tables. Cards as an intervention were analyzed separately in Section 4.5. Here, the question is simply: whose contributed items were more likely to turn into trades?

4.6.1 Age and trading outcomes for real goods

Older students contributed 9 real goods, of which 3 were sold and 6 remained unsold, implying a sale rate of about 33%. Younger students contributed 30 real goods, of which 14 were sold and 16 remained unsold, for a sale rate of roughly 47%. Thus, younger students not only supplied the majority of the inventory but also had a higher probability that their items would be traded at least once. Given the small sample, these differences should be treated as descriptive rather than statistically tested, but they are consistent with the qualitative impression that younger students were more actively engaged with the marketplace.

Table 4.2.

Real goods by age group and sale outcome (non-card items only)

Age group	Not sold	Sold	Total	% sold
Older	6	3	9	33%
Younger	16	14	30	47%
Total	22	17	39	

4.6.2 Gender and trading outcomes for real goods

Female students contributed 13 real goods, of which 4 were sold and 9 remained unsold, corresponding to a sale rate of about 31%. Male students contributed 26 real goods, with 13 sold and 13 unsold, for a sale rate of 50%. As with age, male students were more likely both to contribute items and to see those items converted into actual trades.

Table 4.3.

Real goods by gender and sale outcome (non-card items only)

Gender	Not sold	Sold	Total	% sold
Female	9	4	13	31%
Male	13	13	26	50%
Total	22	17	39	

4.6.3 Inventory composition by gender: books vs. non-books

To assess whether gender differences in sale rates reflect differences in the composition of contributed goods, Table 4.4 presents a 2×2 breakdown of real goods (non-card items only) into books and non-books (toys and other items).

Table 4.4.

Real goods by gender and item type (non-card items only)

Gender	Books	Non-books	Total
Female	7	6	13
Male	13	13	26
Total	20	19	39

Both female and male students contributed a mix of books and non-books. For girls, 7 of 13 items (about 54%) were books; for boys, 13 of 26 items (50%) were books. In other words, the share of books in each group's contributions is very similar. This suggests that the lower sale rates observed for female-supplied items (Section 4.6.2) are not simply an artefact of girls bringing "only books" while boys bring toys. Instead, both genders supplied roughly half books and half non-books; the stronger trading performance of male-supplied items reflects differences in overall engagement and demand, not a fundamentally different composition of what was brought to the shop.

4.6.4 Inventory composition by age: books vs. non-books

A similar check can be done for age groups. Table 4.6 shows the distribution of real goods by Younger vs Older students and Books vs Non-books.

Table 4.6.

Real goods by age group and item type (non-card items only)

Age group	Books	Non-books	Total
Older	4	5	9
Younger	16	14	30
Total	20	19	39

Older students contributed 9 real goods (4 books and 5 non-books), while younger students contributed 30 real goods (16 books and 14 non-books). As with gender, both age groups brought a

mix: books account for roughly 44% of older students' items (4/9) and 53% of younger students' items (16/30). Thus, the higher sale rate for younger students' items (Section 4.6.1) is not merely the result of older students only contributing books. Younger students contributed more of everything, and their items were more likely to sell, while older students participated less and saw a larger share of their mixed inventory remain unsold.

4.6.5 Linking real-goods heterogeneity to the intervention

Taken together, these patterns indicate that the real-goods market was driven mainly by younger male students, whose items sold more frequently than those of older students and girls. This concentration becomes even clearer when viewed alongside the intervention results in Section 4.5: all 25 card purchases were made by younger students, and 24 of 25 by boys. In other words, the same subgroup that dominated successful trades in real goods also made almost exclusive use of the collectible cards. This could be due to the fact, that contributing goods made them earn more tokens, having the ability to spend it.

From a distributional perspective, the school currency system therefore did not generate a broad, even marketplace. Instead, it created a small, relatively active core whose items sold more often and who used the intervention to spend part of their token balances, while other students participated less and had a larger share of their contributed goods left unsold.

4.7 Synthesis of empirical findings

Sections 4.1 to 4.6 point to a market that was able to operate, but only while there was something participants actually wanted to buy. In practice, the token behaved less like general money and more like a claim on a limited shop assortment. Once the most attractive items were sold early on, the remaining goods were much harder to move and token balances stopped circulating. This is why the main issue appears to be circulation rather than an outright shortage of tokens. Demand was selective, and overall activity depended strongly on the composition of the shop's inventory.

This also helps explain why unused value accumulated in two places at the same time. Part of it remained in student balances that were never spent, while another part was tied up in unsold inventory that the shop had already paid for. Paying tokens for incoming goods without strong demand screening effectively turned the shop into an inventory sink, since some of the issued currency became locked into items that did not convert back into purchases.

The timing patterns are consistent with this interpretation. Even when items eventually sold, transactions often occurred with noticeable delays. This points to a low-activity environment in which trade depended on occasional matches rather than a steady flow of exchange.

The card intervention is best understood as a spending outlet rather than as regular inventory. Cards enter the data only when they are purchased, so they represent realized spending rather than stock waiting to be sold. They therefore provided a simple way for existing balances to turn into recorded transactions, without addressing the underlying issue that the real-goods menu could weaken and leave both goods and balances idle.

5. IMPLICATIONS, LIMITATIONS, AND FUTURE RESEARCH

5.1. Summary of key findings

Our experiment set out to explore whether a non-redeemable, school-specific currency could mobilize idle household goods for trade in a thin, closed environment and how design choices shape circulation.

The empirical evidence indicates that the system operated, but in a muted and uneven manner. While the shop successfully facilitated exchanges and generated spending activity, a substantial share of both inventory and issued currency failed to circulate. Much of the issued currency was either absorbed only after the introduction of an auxiliary spending channel or remained idle, alongside unsold goods. Overall, the observed outcomes are consistent with a low-velocity equilibrium characterized by limited participation and a tendency toward hoarding.

On the theoretical side, the mathematical chapter formalized a “hoarding channel” under Cobb–Douglas preferences with a third good, money, entering utility. When money receives a positive utility share, lower expenditure on real goods and reduced trading volume is expected. Even though the Edgeworth-box set of efficient allocations in goods space is unchanged. This formalization provides a clean benchmark for interpreting the empirical pattern of low trade volume despite an apparent abundance of idle goods and tokens.

Together, these elements draw a coherent picture. The experiment generated observable exchange activity, showing that the system could support real transactions among participants. Even with limited participation and weak incentives, the marketplace enabled interaction between buyers and sellers, leading to completed trades and some circulation of the school currency. At the same time, performance was modest relative to the effort and resources required to set up and operate the system. The low volume of exchange points to significant implementation frictions, suggesting that the central issue is not whether such a mechanism can function at all, but whether it can operate efficiently at scale. From this perspective, the experiment is best read as a feasibility exercise that reveals design bottlenecks, rather than as evidence of steady-state performance.

Parallel currencies operate as context-specific institutional arrangements rather than standardized policy tools. Their performance depends on local norms, the composition of participants, trust relations, and the constraints of the community in which they are embedded. Designs that

function in one environment therefore cannot be transplanted elsewhere with any expectation of producing similar outcomes.

This limits the usefulness of searching for a single “optimal” model of parallel currency design. Each implementation instead occupies a different point within a broader design space, with outcomes shaped by the interaction between institutional rules and local conditions. Viewed in this way, the present experiment does not aim to offer a transferable solution. Its contribution lies in making visible the constraints and trade-offs that condition feasibility in a specific setting.

5.2 Parallel currencies, leakage, and acceptance density

The literature on parallel and community currencies highlights two broad conditions under which such systems matter: when they create internal liquidity that cannot easily leak into national money and when they constrain spending to selected networks through time, geographical or merchant rules. The school currency adopted the first logic: non-redeemability, while also restricting acceptance to a single on-campus shop.

From a leakage perspective, the design worked exactly as intended. Unlike redeemable currencies, where value can quickly exit back into national currency, tokens in the school system could not be cashed out and were worthless outside the experiment. The 46-unit net outflow that remained in students’ hands at the end of the project is therefore best interpreted not as leakage but as trapped liquidity: purchasing power that never managed to find attractive uses within the network. In input–output terms, the experiment successfully prevented monetary outflows but did not generate additional internal rounds of spending.

The results also align with work on acceptance density and two-sided payment networks. Theoretical and empirical studies emphasize that tokens create meaningful circulation only when a critical mass of merchants and consumers coordinate on acceptance; otherwise, the option value of waiting dominates and balances sit idle. In the school context, acceptance was extremely thin: there was effectively one merchant (the shop), a small and quickly exhausted menu of desirable items, and no peer-to-peer trading. Under these conditions, it is unsurprising that students accumulated tokens but spent cautiously. The system achieved the “no-leakage” goal of a closed token yet failed to reach the self-sustaining circulation observed in more developed mutual-credit networks.

This contrast is particularly stark when viewed against municipal voucher programs, where redeemability and wide merchant networks produce short-lived but sizeable spending spikes. Our design deliberately traded off ease of use and breadth of acceptance against tight non-redeemability. The experiment suggests that, in a very thin micro-market, closing the cash-out channel is not enough; without variety and acceptance density, the main effect of a non-redeemable token may be to warehouse purchasing power rather than mobilize it.

5.3 Hoardability and the behavioral value of waiting

The mathematical framework clarifies how a positive utility of money per se can reduce trade volume even when underlying preferences over goods remain unchanged. With Cobb–Douglas preferences in three goods, granting money a utility share γ scales down the budget share for real goods by $(1-\gamma)$. Agents optimally allocate more of their income to money balances, lowering equilibrium expenditure on A and B and compressing the volume of goods traded.

The school experiment also illustrates how hoarding interacts with a finite horizon. The end of the project was clearly signaled in advance, and after the final shop session tokens had no further use: effectively a 100% demurrage shock on the last day. In that setting, a purely “high- γ ” story would predict that students run down their balances as the horizon approaches. We did observe that some students, particularly the already-active younger boys, responded exactly this way by spending their remaining tokens on collectible cards in the final phase.

However, a substantial share of tokens were simply allowed to expire unused. This suggests that, even though by design nearing the end of the experiment, the utility of hoarding tokens was actually close to zero, but the perceived utility of the remaining inventory was even lower. What looked like continued hoarding is therefore better interpreted as refusal to trade at unfavorable terms: for some students, the marginal value of a last-minute purchase (mostly of unwanted books) was below the marginal value of keeping tokens as a neutral “souvenir” or doing nothing at all. In model terms, the empirical pattern reflects not only the hoardability parameter γ but also the very low subjective value of the goods still on offer in a thin market with poor inventory and limited acceptance.

Beyond the formal γ channel, the literature on behavioral motives for holding liquidity helps interpret the observed patterns. Students faced uncertainty about future inventory and restocks, which strengthens buffer-stock motives and the option value of waiting. Anticipation of “something better

later” is visible in the data: toys and attractive non-book items sold quickly, while the remaining stock of books provoked little spending despite horizon pressure as the project approached its end. Mental accounting and reference points likely reinforced this inertia, some students informally reported wanting to “keep at least X tokens” or save for a vague “big item,” even when no such item was actually on offer.

In this sense, the empirical case matches the model’s insight: once money provides utility as a store of option value, competitive allocations tilt away from real-goods consumption. The Edgeworth-box prism still contains efficient allocations that exhaust goods, but the equilibrium selected by money-loving agents need not reach them.

The collectible cards intervention confirms and refines this picture. Cards did not change γ directly (as they did not make money less attractive as such) but they created a new, simple way to convert balances into immediate enjoyment. For the active subgroup of younger boys, the perceived marginal utility of spending one token on a card exceeded the marginal utility of holding it. This shifted their private γ downward in practice, raising their individual velocity. Yet for less engaged students or those with few tokens, γ remained high and spending limited, so the intervention’s effect on system-wide trade remained partial.

5.4 Market thinness, inventory composition and distributional patterns

Another key lesson concerns market thinness and the composition of goods. The literature on small-market leakage emphasizes that limited assortment can rationally discourage participation when available goods no longer match preferences. Our results replicate this dynamic in miniature: once the small stock of toys and other attractive items was exhausted, the shop’s inventory became heavily skewed toward books that many students did not value. The accumulation of unsold inventory and the declining pace of trades highlight how quickly a thin market can stall when remaining goods fail to sustain demand.

The distribution of outcomes reinforces this interpretation. A small subset of participants accounted for a disproportionate share of item contributions, successful sales, and use of the card intervention. These students supplied more goods across categories, achieved higher sale rates, and were responsible for nearly all card purchases. Other students contributed broadly similar mixes of books and non-book items but participated less actively in trading, which translated into lower sale

rates and minimal engagement with the intervention. This pattern indicates that unequal outcomes were not driven solely by differences in inventory composition, but by variation in participation intensity. An active core treated the system as a space for repeated interaction, while others remained largely passive.

From a policy design perspective, this concentration is consequential. Parallel currencies are often presented as inclusive instruments for local development, yet in practice their benefits may accrue primarily to already-engaged participants. In this experiment, the interaction between token design and inventory dynamics generated a small group of frequent traders and collectors rather than broad-based participation. This mirrors findings from evaluations of local currency initiatives that report high visibility but limited reach beyond early adopters.

The card intervention appears to have amplified this dynamic. By design, the cards drew on symbols and aesthetics that resonated strongly with a subset of participants and functioned as an effective nudge for those already active in trading. At the same time, they did little to attract students who had not previously engaged with the system. Future implementations are therefore likely to face a trade-off between depth, meaning intensified engagement among a motivated core, and breadth, meaning wider participation and more even acceptance across the community.

5.5 Implications for local-currency design and for school-based experiments

The combination of theoretical and empirical results yields several practical implications. First, in very small economies with limited variety and one main acceptance point, a non-redeemable token can easily become a hoarding device rather than a circulation engine. The school experiment shows that even when leakage is mechanically zero, velocity can remain low because inventories are unattractive and acceptance thin. Designers of local currencies in small communities should therefore treat assortment and acceptance density as first-order design levers, on par with redeemability and demurrage.

Second, soft spending incentives such as collectibles or loyalty-style rewards can be effective but inherently limited. In this experiment, the introduction of collectible cards drew back a meaningful share of previously idle tokens and accounted for a large fraction of observed transactions. At the same time, this recovery was driven mainly by increased spending among participants who were already active, rather than by a broader expansion of supply. The intervention did not lead to a visible

increase in new goods being offered for sale, indicating that such incentives are better at mobilizing existing balances than at inducing additional participation on the supply side.

For policy design, this points to a clear trade-off. Targeted and playful spending outlets can efficiently absorb idle balances within specific segments of a community, but they do not automatically generate wider engagement or more inclusive market activity. Without complementary mechanisms that encourage contribution and acceptance, these tools may end up reinforcing existing participation gaps rather than narrowing them.

Finally, this small case study hints at why it can be useful to think explicitly about hoardability when discussing parallel currencies. Some debates about leakage and multipliers focus on whether value escapes into national money, but in our setting, there was no cash-out option and value still ended up sitting idle because the acceptance network was thin and the remaining inventory was unattractive. While the experiment is too small to support strong policy claims, it does illustrate on a micro scale that, alongside formal design features, designers of school-based or local schemes may want to consider both why people might prefer to hold tokens and how many genuinely appealing ways there are to spend them.

5.6 Limitations and avenues for future research

Methodologically, the study comes from a single school and a single implementation of the shop, and the dataset consists of 64 items that students happened to bring in. The size and composition of the sample were not controlled in advance but emerged endogenously through participation. This limits statistical power and any claim that the observations are representative of a wider population. The experiment also lacked a control group or randomized interventions: the collectible cards were introduced adaptively in response to observed hoarding rather than as a pre-planned treatment. For these reasons, any causal interpretation of the findings has to remain modest and exploratory.

Measurement is also constrained. No survey data were collected on students' motivations, perceptions of value or reasons for hoarding versus spending, so behavioral interpretations rely on informal observations and theoretical priors rather than direct elicitation. The dataset does not capture peer-to-peer trades that might have occurred outside the shop, nor does it track how many tokens each student held at different points in time.

On the theoretical side, the mathematical model abstracts from many features that may matter in practice. The Cobb–Douglas framework with a single hoardable money share captures the key

hoarding channel but ignores heterogeneity in preferences, network effects, social influence, and non-Cobb–Douglas behaviors such as thresholds or satiation. Behavioral sections in the literature review highlight richer mechanisms: mental accounting, goal thresholds, scarcity mindset, that could be incorporated into future models.

These limitations point to several avenues for future work:

Scaling up and replication: running similar experiments in multiple schools or youth clubs, with larger inventories and longer horizons, would allow for more robust statistical analysis and exploration of cross-context variation.

Design comparison: Future studies could systematically vary key design levers (redeemability, demurrage, acceptance density, initial endowments) and compare outcomes across treatment arms rather than relying on a single adaptive intervention.

Richer data collection: Combining transaction logs with surveys, simple experiments (e.g. willingness-to-spend tasks) and perhaps digital wallets would enable direct estimation of γ -like parameters, spending propensities and behavioral types.

Model extensions: The formal framework could be extended to include heterogeneous γ , explicit two-sided network effects, and stochastic inventory, bridging the gap between search-theoretic monetary models and the behavioral channels summarized in the literature review.

REFERENCES:

- Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics*, 37(3), 668–691. https://www.researchgate.net/publication/24049369_Compensation_in_Two-Sided_Markets
- Baines, R., de Freitas, C., Ferreira, N., & Land, M. (2020). Community exchange and time currencies: A systematic and integrative review. *Systematic Reviews*, 9(1), 1–18. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7093815/>
- Blanc, J. (2011). Classifying “CCs”: Community, complementary and local currencies’ types and generations. *International Journal of Community Currency Research*, 15(D), 4–10. <https://ijccr.net/wp-content/uploads/2012/05/ijccr-2011-special-issue-02-blanc.pdf>
- Borgatti, S. P., Mehra, A., Brass, D. J., & Labianca, G. (2009). Network analysis in the social sciences. *Science*, 323(5916), 892–895. https://www.researchgate.net/publication/24004962_Network_Analysis_in_the_Social_Sciences
- Caillaud, B., & Jullien, B. (2003). Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics*, 34(2), 309–328. https://www.researchgate.net/publication/24049216_Chicken_Egg_Competition_Among_Intermediation_Service_Providers
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica*, 59(5), 1221–1248. https://www.princeton.edu/~deaton/downloads/Saving_and_Liquidity_Constraints.pdf
- Duffy, J., & Puzzello, D. (2014). Gift exchange versus monetary exchange: Theory and evidence. *American Economic Review*, 104(6), 1735–1776. <https://www.aeaweb.org/articles?id=10.1257/aer.104.6.1735>
- Easley, D., & Kleinberg, J. (2010). *Networks, crowds, and markets*. Cambridge University Press <https://www.cs.cornell.edu/home/kleinber/networks-book/networks-book.pdf>
- Fleischman, T., & Dini, P. (2020). Liquidity-saving through obligation-clearing and mutual credit. *Journal of Risk and Financial Management*, 13(12), 295. <https://ideas.repec.org/a/gam/jjrfmx/v13y2020i12p295-d452515.html>
- Geier, A. B., Rozin, P., & Doros, G. (2006). Unit bias: A new heuristic that helps explain the effect of portion size on food intake. *Psychological Science*, 17(6), 521–525. <https://journals.sagepub.com/doi/10.1111/j.1467-9280.2006.01738.x>
- Gelleri, C. (2009). Chiemgauer Regiomoney: Theory and practice of a local currency. *International Journal of Community Currency Research*, 13, 61–75. <https://core.ac.uk/download/pdf/12174396.pdf>

- Glaeser, E. L., Kolko, J., & Saiz, A. (2000). Consumer city. *Journal of Economic Geography*, 1(1), 27–50. https://scholar.harvard.edu/files/glaeser/files/consumer_city.pdf
- Hayashi, K. (2021). Rethinking the significance of regional currencies: The case of the Chiemgauer. *International Journal of Community Currency Research*, 25, 35–47. https://ijccr.net/wp-content/uploads/2021/03/96_106-ijccr-2021_1-hayashi.pdf
- Hoch, S. J., Bradlow, E. T., & Wansink, B. (1999). The variety of an assortment. *Marketing Science*, 18(4), 527–546. <https://pubsonline.informs.org/doi/10.1287/mksc.18.4.527>
- Hua, C.-C. (2022). Evaluating economic benefits of the revitalization vouchers: Evidence from Taiwan's Triple Stimulus Vouchers. *Pacific Economic Review*, 27(4), 319–339. <https://ideas.repec.org/a/bla/pacecr/v27y2022i4p319-339.html>
- Hughes, D. W. (2018). *A primer in economic multipliers and impact analysis*. University of Tennessee Extension. <https://utia.tennessee.edu/publications/wp-content/uploads/sites/269/2023/10/W644.pdf>
- Jackson, M. O. (2010). *Social and economic networks*. Princeton University Press. https://www.researchgate.net/publication/331307250_Social_and_Economic_Networks
- Johnson, S., & Harvey-Wilson, S. (2018). The Bristol Pound: An evaluation. *Local currency adoption and use: Insights from a realist evaluation of the Bristol Pound (Bath Papers in International Development and Wellbeing No. 56)*. University of Bath. <https://www.econstor.eu/bitstream/10419/223124/1/1040937446.pdf>
- Kim, M., & Oh, Y. H. (2021). The impact of COVID-19 regional cash subsidies on the sales of local businesses in South Korea. *KDI Journal of Economic Policy*, 43(2), 103–123. https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID3870948_code2536162.pdf?abstractid=3870948&mirid=1
- Kiyotaki, N., & Wright, R. (1993). A search-theoretic approach to monetary economics. *American Economic Review*, 83(1), 63–77. [https://homepage.ntu.edu.tw/~yitingli/file/macro_and_money/Kiyotaki%2C N. and R. Wright %281993%29.pdf](https://homepage.ntu.edu.tw/~yitingli/file/macro_and_money/Kiyotaki%2C%20N.and%20R.Wright%281993%29.pdf)
- Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The goal-gradient hypothesis resurrected. *Journal of Marketing Research*, 43(1), 39–58. <https://journals.sagepub.com/doi/abs/10.1509/jmkr.43.1.39>
- IMPLAN. (2023). *Input–Output & Social Accounting Matrix structure*. IMPLAN Group, LLC.
- Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy*, 99(3), 483–499. <https://www.jstor.org/stable/2937739>
- Lagos, R., & Wright, R. (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113(3), 463–484. https://qed.econ.queensu.ca/pub/students/phds/liuqian/MRG/Fall_2007/Lagos_Wright2005.pdf

- LM3/NEF. (n.d.). *Local Multiplier 3 (LM3) methodology overview*. NEF Consulting. <https://nefconsulting.com/wp-content/uploads/2017/05/TheMoneyTrail.pdf>
- Marshall, A. P. (2018). The Bristol Pound: A tool for localisation? *Ecological Economics*, 146, 273–281. https://repository.essex.ac.uk/30114/1/BP_Paper_finalpdf.pdf
- Matti, J., & Zhou, Y. (2022). Money is money: The economic impact of BerkShares. *Ecological Economics*, 192, 107255. <https://ideas.repec.org/a/eee/ecolec/v192y2022ics0921800921003141.html>
- Newman, M. E. J. (2010). *Networks: An introduction*. Oxford University Press. <https://academic.oup.com/book/27303>
- Nosal, E., & Rocheteau, G. (2011). *Money, payments, and liquidity*. MIT Press. <https://mitpress.mit.edu/9780262533270/money-payments-and-liquidity/>
- Nunes, J. C., & Drèze, X. (2006). The endowed progress effect: How artificial advancement increases effort. *Journal of Consumer Research*, 32(4), 504–512. <https://academic.oup.com/jcr/article-abstract/32/4/504/1787425>
- Opsahl, T., Agneessens, F., & Skvoretz, J. (2010). Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3), 245–251. <https://www.sciencedirect.com/science/article/abs/pii/S0378873310000183>
- Prelec, D., & Loewenstein, G. (1998). The red and the black: Mental accounting of savings and debt. *Marketing Science*, 17(1), 4–28. <https://doi.org/10.1287/mksc.17.1.4>
- Parker, G. G., & Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management Science*, 51(10), 1494–1504. https://www.researchgate.net/publication/220534822_Two-Sided_Network_Effects_A_Theory_of_Information_Product_Design
- Reyns, A., van Lerven, F., & others. (2025). How mutual credit systems bolster resilience of small firms. *Journal of Economic Behavior & Organization* (forthcoming). <https://www.sciencedirect.com/science/article/abs/pii/S016726812400475X>
- Rochet, J.-C., & Tirole, J. (2005). Two-sided markets: A progress report. *International Journal of Industrial Organization*, 23(7–8), 1089–1112. https://www.tse-fr.eu/sites/default/files/medias/doc/by/rochet/rochet_tirole.pdf
- Rötheli, T. F. (2009). *The superiority of monetary over barter exchange: Experimental results and policy implications*. https://www.researchgate.net/publication/227377191_The_superiority_of_monetary_over_barter_exchange_Experimental_results_and_policy_implications
- Rysman, M. (2007). An empirical analysis of payment card usage. *Journal of Industrial Economics*, 55(1), 1–36. <https://doi.org/10.1111/j.1467-6451.2007.00301.x>

Schiff, N. (2014). Cities and product variety: Evidence from restaurants. *Journal of Economic Geography*, 15(6), 1085–1123. <https://doi.org/10.1093/jeg/lbu040>

Seyfang, G., & Longhurst, N. (2013). Growing green money? Mapping community currencies for sustainable development. *Ecological Economics*, 86, 65–77. <https://ideas.repec.org/a/eee/ecolec/v86y2013icp65-77.html>

Soman, D. (2001). Effects of payment mechanism on spending behavior: The role of rehearsal and immediacy of payments. *Journal of Consumer Research*, 27(4), 460–474. <https://doi.org/10.1086/319621>

Srivastava, J. (2009). The denomination effect. *Journal of Consumer Research*, 36(4), 701–713. <https://doi.org/10.1086/599222>

Stodder, J., & Lietaer, B. (2016). The macro-stability of Swiss WIR-Bank credits: Balance, velocity, and leverage. *Comparative Economic Studies*, 58(4), 570–605. https://ideas.repec.org/a/pal/compes/v58y2016i4d10.1057_s41294-016-0001-5.html

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, 12(3), 183–206. [https://doi.org/10.1002/\(SICI\)1099-0771\(199909\)12:3](https://doi.org/10.1002/(SICI)1099-0771(199909)12:3)

SUMMARY

65 pages, 7 tables, 3 charts, 43 references.

This thesis examines whether a closed-loop, non-redeemable “school currency” can mobilize idle household goods for exchange in a thin market, and which constraints most strongly limit circulation in practice. The work aims to evaluate the feasibility of such a mechanism in a school setting and to identify the main drivers of participation, hoarding, and low velocity. The objectives include synthesizing key design levers from the parallel-currency literature, formalizing a benchmark mechanism for hoarding, implementing an embedded pilot with clear ethical boundaries, and analyzing item-level outcomes before and after a spending nudge.

Methodologically, the thesis combines a literature-based mechanism review of complementary/parallel currency designs and circulation constraints, a conceptual/formal benchmark that models “hoarding” as a positive utility share for money under Cobb–Douglas preferences, and a descriptive case study based on transaction logs. Empirically, an eight-week pilot was run in a private school in Vilnius with an accessible pool of ~115 students (grades 5-11). Transactions were recorded as an item-level spreadsheet log; after cleaning, the analytical sample contains 64 entries.

Results show that the marketplace functioned but remained low-volume and uneven. For real goods, 39 items were listed and 17 sold ($\approx 44\%$ sale rate), with strong category differences: books sold poorly (3/20; 15%), while toys were highly liquid (12/13; 92%). Trades were often delayed (median time-to-sale ≈ 8 days). From an accounting perspective, the shop issued 123 units to acquire real goods but recovered only 52 units through resales ($\approx 42\%$), leaving 71 units effectively tied up in unsold inventory.

The main conclusion is that a non-redeemable token can enable real exchange yet still settle into a low-velocity equilibrium when acceptance is thin and the amount of desirable goods is low. In this setting, non-redeemability prevented cash-out leakage, but a substantial share of value became trapped as idle balances and unsold stock. Soft spending incentives can absorb hoarded balances within an active subgroup, but do not automatically solve inventory scarcity or participation gaps. Limitations include the single-school setting, modest and endogenous inventory, adaptive (non-randomized) intervention timing, and anonymity constraints that prevent individual-level tracking.

The results have not been published to date, however, it is suitable for development into a short case-study or methods-oriented publication on circulation constraints in micro-scale currency pilots.

SANTRAUKA

65 puslapiai, 7 lentelės, 3 paveikslai, 43 šaltiniai.

Šiame darbe nagrinėjama, ar į eurus nekeičiama mokyklinė valiuta gali paskatinti, namuose nenaudojamų žemos vertės daiktų mainus ir kokie veiksniai daro daugiausia įtakos tokios valiutos apyvartai. Darbo problema kyla iš to, kad net ir sukūrus atsiskaitymo priemonę, apyvarta gali išlikti menka, jei dalyviai linkę kaupti alternatyvią valiutą arba jei „parduotuvės“ pasiūla tampa nepatraukli. Tyrimo klausimas: ar mokyklinė, nekeičiama valiuta įgalina mainus, ir kas lemia apyvartą.

Darbo tikslas - įvertinti paralelinės mokyklinės valiutos galimybes mobilizuoti nenaudojamus daiktus mainams ir identifikuoti mechanizmus, lemiančius apyvartą bei dalyvių įsitraukimą. Tam išskelti uždaviniai: atlikti literatūros apie paralelines ir papildomas valiutas analizę, įvertinti pinigus kaip prekę iš matematinės perspektyvos, atlikti pilotinį eksperimentą, išanalizuoti vykusius sandorius, aprašant pasiūlos struktūrą, pardavimus, dalyvių įsitraukimą ir reaktyviai vykdytų intervencijų poveikį.

Literatūros analizėje aptariamos ir klasifikuojamos didžiausio pasisekimo sulaukusios paralelinės valiutos, nagrinėjami jų panašumai, skirtumai, veiksniai darantys įtaką projektų pasisekimui. Taip pat nagrinėjamas rinkos dalyvių noras kaupti pinigus, jo priežastys ir pasekmės. Praktinė dalis atlikta Vilniaus privačioje mokykloje. Tyrimo metu buvo išleista mokyklos valiuta ir atidaryta parduotuvė, kurioje vaikai galėjo mainyti nenaudojamus daiktus į alternatyvią valiutą ar įsigyti kitų mokinių atneštus daiktus. Jame galėjo dalyvauti 115 mokinių iš 5–11 klasių. Pilotinis tyrimas truko aštuonias savaites. Rinkti duomenys apie prekių kainą (vietinė valiuta), kategoriją, pardavimo laiką. Imtį sudaro 64 įrašai.

Rezultatai rodo, kad parduotuvė veikė, tačiau buvo maža ir netolygi. Iš 39 prekių, kurias atnešė mokiniai parduota 17 (apie 44%); knygos buvo mažiausiai likvidžios (3/20; 15%), tuo tarpu žaislai - beveik visiškai išparduoti (12/13; 92%). Sandoriai vyko gan lėtai (pardavimo laiko mediana 8 dienos). Apskaitiniu požiūriu parduotuvė pirkdama prekes išmokėjo 123vnt. valiutos, o iš realių prekių pardavimų grįžo 52vnt. todėl 71vnt. liko „įstrigęs“ neparduotose prekėse. Intervencija sugeneravo 25 ir gražino papildomus 25vnt. valiutos, sumažindama parduotuvės nuostolį nuo 71 iki 46vnt.

Pagrindinė išvada – bandant įvesti alternatyvią valiutą neužtenka apriboti galimybę iškeisti ją į eurus. Jei neužtikrinama, jog valiuta turės vertę, t.y. rinkos dalyviai galės ją išsikeisti į norimas prekes ar paslaugas dėl mažo asortimento ir neprognozuojamų inventoriaus papildymų, dalyviai bus linkę ją kaupti taip pat, ar net labiau nei įprastas valiutas.

Darbo rezultatai šiuo metu nėra publikuoti, tačiau pilotinis tyrimas gali būti naudojamas kaip pagrindas tolimesniems paralelinių valiutų eksperimentams panašiose socio-ekonominėse aplinkose.

APPENDIXES:

Appendix 1.

Evidence-anchored comparison of archetypes

Dimension	Redeemable city currencies (Bristol, BerkShares)	Non-redeemable mutual credit (WIR, Sardex)
Convertibility	Yes (par to fiat) → MPR > 0 by design	No cash-out → MPR = 0
Implication for MPC	Sustained MPC low unless strict constraints bind	MPC feasible high if acceptance curated and supply credible
Observed outcomes	High visibility; durable localization limited; macro effects often null (BerkShares)	Countercyclical use; recirculation via clearing; stabilization among SMEs
What props it up	Discounts, publicity, reassurance of convertibility	Merchant curation, credit lines, multilateral netting, risk management
Typical failure mode	Leakage via MPR; reversion to base money post-program	If acceptance/supply thin → MPS rises, velocity stalls

Appendix 2

Experiment currency design



Appendix 3

Design of collectible cards issues to stimulate trade



Appendix 4

Dataset:

Category	Age	Gender	Price Bought	Time sold (D)	Price Sold	Profit
Toy	Older	F	2	0	2	0
Toy	Older	F	2	0	2	0
Book	Older	F	2	not sold		-2
Book	Older	F	2	not sold		-2
Book	Older	F	6	not sold		-6
Book	Older	F	3	not sold		-3
Toy	Older	F	3	not sold		-3
Other	Older	F	3	not sold		-3
Book	Younger	M	3	not sold		-3
Book	Younger	M	3	2	3	0
Book	Younger	M	2	not sold		-2
Book	Younger	M	2	not sold		-2
Book	Younger	M	2	7	2	0
Toy	Younger	M	1	3	1	0
Other	Older	M	5	0	6	1
Book	Younger	M	4	not sold		-4
Book	Younger	M	3	not sold		-3
Book	Younger	M	3	not sold		-3
Other	Younger	M	2	not sold		-2
Book	Younger	M	2	not sold		-2
Toy	Younger	M	7	8	8	1
Book	Younger	M	5	not sold		-5
Book	Younger	F	3	8	3	0
Book	Younger	F	3	not sold		-3
Book	Younger	F	3	not sold		-3
Book	Younger	M	4	not sold		-4
Book	Younger	M	6	not sold		-6
Book	Younger	M	3	not sold		-3
Toy	Younger	M	3	0	3	0
Toy	Younger	M	1	11	1	0
Toy	Younger	M	1	52	1	0

Other	Younger	M	1	not sold			-1
Toy	Younger	M	3	11	3	0	
Toy	Younger	M	3	11	3	0	
Toy	Younger	M	3	11	3	0	
Toy	Younger	M	3	11	3	0	
Other	Younger	F	8	43	5	-3	
Toy	Younger	M	3	0	3	0	
Other	Younger	F	5	not sold			-5
Our Cards	Younger	F		2	1	1	
Our Cards	Younger	M		12	1	1	
Our Cards	Younger	M		11	1	1	
Our Cards	Younger	M		10	1	1	
Our Cards	Younger	M		9	1	1	
Our Cards	Younger	M		8	1	1	
Our Cards	Younger	M		7	1	1	
Our Cards	Younger	M		6	1	1	
Our Cards	Younger	M		5	1	1	
Our Cards	Younger	M		4	1	1	
Our Cards	Younger	M		3	1	1	
Our Cards	Younger	M		2	1	1	
Our Cards	Younger	M		15	1	1	
Our Cards	Younger	M		14	1	1	
Our Cards	Younger	M		13	1	1	
Our Cards	Younger	M		12	1	1	
Our Cards	Younger	M		11	1	1	
Our Cards	Younger	M		10	1	1	
Our Cards	Younger	M		9	1	1	
Our Cards	Younger	M		8	1	1	
Our Cards	Younger	M		7	1	1	
Our Cards	Younger	M		6	1	1	
Our Cards	Younger	M		5	1	1	
Our Cards	Younger	M		4	1	1	
Our Cards	Younger	M		3	1	1	