Payment-Chain Crises*

Saki Bigio[†] UCLA and NBER

Esteban Méndez[‡] Central Bank of Costa Rica

Diana Van Patten[§] Yale University and NBER

October 2026

Abstract

This paper introduces an endogenous network of payment chains into a business-cycle model. Some production orders are paid upfront, while others are settled only after the buyer receives payment from its own customers. When production or fulfillment waits for settlement, liquidity shortages can propagate through the payment network, leave resources idle, and cause a reduction in measured productivity. We bring the model to Costa Rican data linking interbank business payments to the universe of electronic invoices, connecting contractual terms, payment timing, and delivery timing. We identify payment chains: firms that pay suppliers only after being paid by customers. Exploiting an unexpected freeze of buyers' funds, we find evidence of TFP declines associated with payment delays. We then embed payment chains into a business-cycle model where firms decide to rely on chained payments. Liquidity shocks delay transactions, feed back into aggregate productivity, and generate hysteresis and pecuniary externalities.

Keywords: Payments, networks, business cycles.

JEL Classification: E32, E42, G01, O16

*We thank Fernando Alvarez, David Argente, Luigi Bocola, Ricardo Lagos, Jennifer La'O, Ezra Oberfield, Michael Peters, Guillaume Rocheteau, Gary Richardson, Pierre-Olivier Weill, and Yuliy Sannikov as well as seminar participants at UC Irvine, UCLA, Princeton, SARG, NBER IFM and MM, CEMLA, and CESifo, for early discussions. Sebastian Merkel, Gadi Barlevy, and Romain Ranciere provided thoughtful discussions. We are indebted to Ken Miyahara and Luis Yepez for their outstanding assistance, multiple discussions, and their creativity in tackling proofs. An earlier version of this work circulated as NBER Working Paper No. 30859; the present paper substantially extends and supersedes that analysis. The views expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Costa Rica (BCCR). All results have been reviewed by the BCCR to ensure no confidential information is disclosed.

[†]Email: sbigio@econ.ucla.edu. Address: 315 Portola Plaza Los Angeles, CA 90095.

[‡] Email: mendezce@bccr.fi.cr. Address: Avenida Central y Primera, Calles 2 y 4, San José 10101.

[§]Email: diana.vanpatten@yale.edu. Address: 165 Whitney Avenue, New Haven, CT 06511.

1 Introduction

A classic narrative has it that during financial crises, productive resources remain idle, not because of technical constraints or insufficient demand, but because agents must wait for others to pay, who in turn are waiting on others. A simple example illustrates the idea. A general contractor is hired to build a commercial property. The contractor relies on a network of subcontractors—electricians, plumbers, and carpenters—who agree to proceed with their work only after receiving partial payments. The contractor also places orders with material suppliers, who customize materials but are unwilling to ship them until payment is completed or, at a minimum, proof of funds is provided. However, the contractor is awaiting payment from a developer, who is awaiting approval of a credit line. During crises, banks may withdraw their credit lines, disrupting the developer's funding sources and, therefore, the entire chain: the contractor cannot pay the electrician or the supplier; the electrician postpones work; the supplier holds off on shipping; and the project stalls. The slowdown is not due to a lack of labor, materials, or technical capacity, but rather to each party waiting to be paid by someone else. This cascading leaves productive human and physical resources idle. National accounts record the outcome as a total factor productivity (TFP) decline.

Despite how natural such payment delays seem, they are rarely featured in macroeconomic models. Production is typically centralized with spot market exchanges and payment frictions abstracted away. While credit constraints appear often in models, the usual assumption is binary: agents are either constrained or unconstrained. These theories explain why resources are misallocated, but not why they sit idle. What is missing is a tractable framework where economic activity stalls simply because agents are waiting for transactions to clear: A firm that cannot pay upfront has to wait for the arrival of payments from its own customers in order to receive inputs. Economic activity can therefore stall not because the firm cannot produce or because demand is insufficient, but because production or delivery awaits settlement.

This paper develops and tests a theory of these payment-chain disruptions. To our knowledge, this is the first paper to model bilateral payment obligations as a macroeconomic network in which the timing of settlement governs the timing of production and thereby shapes TFP. Moreover, we not only propose a new theory but also test this mechanism using economy-wide data that link contractual terms, realized payment timing, delivery timing, and firm outcomes.

The starting point of the paper is a payment-chain production network that distinguishes two objects: spot transactions, which are paid upfront, and chained transac-

tions, which remain unfulfilled until preceding payments clear. Spot transactions encompass any transaction in which goods are delivered immediately: either because the firm has retained funds or has obtained credit from the financial system or directly from its suppliers. Instead, chained transactions encompass trading agreements in which the good is delivered only when the customer settles a payment. We show that, given a set of spot and chained orders, the production network determines how settlement delays are allocated across transactions. This block yields a tractable mapping from micro-level transaction assignments to aggregate TFP via the network: as the share of chained orders rises, settlement delays cascade through longer chains and reduce measured TFP.

We further demonstrate how, when firms make many transactions, this payment-chain block can be aggregated into a fictitious budget constraint where firms face a price for buying goods through spot or chained transactions. In the aggregate budget constraint, chained orders appear as goods purchased at an endogenous price wedge. This aggregation reveals two sources of inefficiency. The first is the static organization of payments: payments may be arranged in unnecessarily long or unequal chains. The second is a pecuniary externality, which plays a central role in the business-cycle model: individual agents do not internalize how their expenditure choices affect the speed of transactions faced by others. This aggregation makes our theory portable to many settings.

We then identify chained transactions and their consequences for TFP using unusually rich administrative data from Costa Rica. We link the universe of interbank business-to-business payments with the electronic invoice issued *for every formal sale* in the economy. For each transaction, the invoice identifies the buyer, the seller, the product, the price, whether payment is immediate or deferred, and, when payment is deferred, the contractual due date. We then link these contractual terms to actual bank payments and firm outcomes. Moreover, direct physical delivery dates are almost never observed in economy-wide administrative data. We make progress on this otherwise unobservable margin by leveraging payments to delivery and transportation providers, yielding an economy-wide proxy for delivery activity. These linked records allow us to study payment obligations and the terms agreed upon at the time of sale, together with their settlement and their consequences for delivery and real economic activity.

The empirical section is organized around four facts that trace the model's mechanism from the creation of a payment obligation, through settlement and delivery, to prices and real activity. First, deferred-payment invoices are common, directly measur-

able, and strongly related to an independent classification of chained behavior based on the timing of incoming and outgoing bank payments. Second, contractual payment deadlines predict product delivery activity: Delivery-service payments rise when deferred invoices come due, and, in retail, shipping payments spike in the week a newly purchased product first appears in the retailer's sales. These two facts establish the ample presence of chained transactions where settlement waits for previous settlements and affects delivery times. Because this form of trade entails a delay in the delivery of goods, it is distinct from trade credit relations in which there are no delays. Third, deferred-payment transactions in which delivery waits until full payment is due carry a substantial price premium, which connects to the price wedge in our aggregation result.

The fourth fact comes from a natural experiment that shows how a financial shock propagates through payment chains into real activity: an unexpected liquidity shock interrupted an expected source of liquidity and immediately reduced the outgoing payments of directly exposed firms. The shock is associated with steeper declines in TFP, sales, and value added among sellers one and two links upstream when the interrupted payment was contractually due sooner, leaving sellers little time to adjust. Complementing this evidence, we show directly exposed firms become more likely to switch from spot to chained behavior after the freeze, and this switch is associated with larger activity declines among both the exposed firms and their upstream suppliers. Together, these results show how a financial shock can propagate through a network of payment obligations and generate declines in real activity beyond the directly exposed firms, motivating our business-cycle analysis.

To study the aggregate implications of this propagation mechanism, we exploit our budget-constraint representation to embed the payment-chain block in a parsimonious business-cycle model with entrepreneurs and creditors. Creditors provide funding and make spot expenditures, whereas entrepreneurs choose how much to borrow and whether to buy inputs through spot or chained orders. Thus, the economy's reliance on payment chains is endogenous. We study how liquidity shocks to entrepreneurs induce a greater reliance on chained orders, leading to greater settlement delays that ultimately depress aggregate measured TFP.

In the business-cycle analysis, this pecuniary externality leads to novel implications. While entrepreneurs internalize the delays they themselves incur, they do not internalize the delays their payment choices impose on others. This feedback generates three distinct regions of the state space: efficient steady states in which all purchases are

made on the spot and TFP is maximal; temporary payment-chain crises in which reliance on chained orders recedes as entrepreneurs deleverage; and persistent crises in which debt overhang weakens incentives to deleverage, leaving the economy persistently reliant on chained orders and with TFP permanently depressed.

This business-cycle setting also yields novel policy insights: A Ramsey planner recognizing payment externalities finds that, during payment-chain crises, creditors spend too little on spot orders, while entrepreneurs spend too much on chained ones. In turn, fiscal policy can have positive multipliers during payment-chain crises—but only if the government makes spot expenditures. If the government spends without paying upfront, its expenditures are detrimental. This gives a new interpretation of fiscal multipliers: in a payment-chain crisis, a government that spends upfront speeds up payments. Thus, how the government spends is what matters.

Literature Review. The paper connects with theories that underscore the sharp declines in aggregate TFP during crises. The link between financial crises and TFP is not at all obvious because financial crises can manifest themselves through labor-market distortions, not productivity. One branch of the literature explains declines in aggregate TFP through increased capital misallocation or capital utilization—see Meza and Quintin (2007), Pratap and Urrutia (2012) or Oberfield (2013). A common finding is that these models can only partially explain TFP declines once disciplined with data. Our paper offers a different mechanism where financial conditions impact TFP through a slowdown of payments and economic activity.

The paper also falls at the crossroads of several areas: the monetary-payments literature, the economic-networks literature, and the literature on aggregate-demand externalities. The issue of how payment instruments affect production is a classic theme.[1] Recent work focuses on how the distribution of funds, and not the quantity of funds, affects production—see Lippi et al. (2015), Rocheteau et al. (2016), and Brunnermeier and Sannikov (2017). As in this literature, the distribution of funding affects allocations, and the main distinction is that we focus on delays in sequential payments.

Sequential payments appear in many other studies. The payments-chain network is inspired by Townsend (1980) or Kiyotaki and Moore (1997a), which can be seen as models where production is distorted within a vertical supply chain.[2] Recent work by Hardy

[1] Lucas and Stokey (1987) analyzes a stochastic cash-in-advance economy; Kiyotaki and Wright (1989) studies trade with indivisible tokens; Lagos and Wright (2005) a model with divisible money and explicit trading arrangements. See also Shi (1997); Lagos et al. (2011); Lagos and Rocheteau (2009); Li et al. (2012); Nosal and Rocheteau (2011); Rocheteau (2011) for many other directions in that area.

[2] Other models of sequential payments include sequences of payments with spatial separation, Freeman (1996) and Green (1999), which study sequential transactions in overlapping-generations environ-

et al. (2022) and Bocola (2022) contrasts external funding against trade credit along the supply chain.[3] Relative to these papers, there are two distinctions. First, we identify a different form of trade from trade credit. Under trade credit, the product is delivered before it is paid for; under a chained order, delivery waits until earlier payments are executed. The empirical literature typically uses receivables and payables as a proxy for trade credit, but these balances can also indicate commitments to pay and deliver. Our data allow us to distinguish situations in which the product is delivered before payment from those in which delivery waits for it. Second, the payments network emerges as an outcome of endogenous expenditure decisions.

With respect to the networks literature, the paper connects with endogenous network formation models. The contribution, however, is modest, as network formation is not strategic. By contrast, in Oberfield (2018), a network is formed through strategic partnerships; in Kopytov et al. (2022) and Elliott et al. (2022) firms form strategic links anticipating possible supply-chain breakdowns. Here, the network is randomly formed, but the distribution of chains is endogenous to financial decisions. As in Elliott et al. (2014), Alvarez and Barlevy (2021) and Taschereau-Dumouchel (2022), there are externalities. In those models, externalities arise when a default sets off further defaults. Here, externalities occur through payment delays.

Finally, the paper connects with models of aggregate demand externalities. An early model of these externalities is Diamond (1982) where, via search, consumption decisions affect output. In most of the literature, demand externalities result from nominal rigidities. There has been recent interest in coupling nominal rigidities with financial constraints—for example, Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2017).[4] The nature of demand externalities here is different. In particular, the type of expenditures by the private or public sector matters: spot orders may stimulate output, but chained orders depress it. The demand externalities arising from payment slowdowns are part of the classic narrative of payment slowdowns.

---

ments, La'O (2015), which studies a circular flow of transactions, or Guerrieri and Lorenzoni (2009), which studies sequential transactions in a Lagos and Wright-type environment. Bigio and La'O (2013) considers the propagation of financial shocks that induce misallocation in a production network.

[3] See also Biais and Gollier (1997) for an earlier model of trade credit, and Petersen and Rajan (1997), Burkart and Ellingsen (2004), and Cuñat (2007) for canonical theories of why suppliers extend credit. This body of work is substantiated by a body of empirical evidence: Boissay and Gropp (2007), Jacobsen (2015), Barrot (2016), and Costello (2020) among others.

[4] Recent papers have further introduced sequential transactions into environments with nominal rigidities—for instance, Woodford (2022) and Guerrieri et al. (2022). In those models, demand externalities occur when agents cut back on any form of expenditure.

2 Payment-Chains and Productivity

This section develops a framework that maps spot and non-spot orders into a payment-chain network. When transaction delays cause delays in transferring goods, utilized time is wasted and chained transactions lead to measured TFP losses. Then, we aggregate the network into a portable budget constraint with externalities.

Bilateral Relations. Consider an economy where production is organized through bilateral trade for inputs. We can think of these inputs as customized construction materials or services, and more generally as specially engineered products. For simplicity, we assume that all input goods are ex-ante identical, before any customization. There are $N \in \mathbb{Z}^+$ such inputs—we work in the space of natural numbers to formalize trade in a network setting. We normalize the total value of intermediate-input sales to one. By assumption, the transaction value and traded quantity in each bilateral trade are identical across orders, even though orders will settle at different instances. Since all N transactions have the same size, each input lot has price $p = 1/N$—price discrimination is ruled out. Inputs are, thus, the numeraire.

The primitives in this section are two types of orders: *spot* and *chained* orders. Spot orders are funded independently of preceding payments in one's in the chain; we refer to them as funded nodes. Funding may come from own resources, the financial system, or trade credit that permits delivery and onward financing without waiting for those payments. From the perspective of our theory, these sources are equivalent; e.g., it does not matter if orders are financed by issuing IOUs to the banking system or to a supplier. The key is that chained orders are contingent on previous payments: the customer finances an order with payment received from another transaction in which he acts as a supplier. There are $N^s \in \mathbb{Z}^+$ spot and $N^x \in \mathbb{Z}^+$ chained orders, with $N = N^s + N^x$ input goods. This "market-clearing condition" will emerge in equilibrium in Section 4.

Each input unit is assigned an identifier, $i \in \mathcal{N} = \{1, 2, \ldots, N\}$. Since the number of inputs sold corresponds to the number of inputs ordered, each order is assigned a unique identifier, $i \in, the index i corresponds to both an input sold and someone's purchase order. We partition \mathcal{N}$ and $\mathcal{N}^x$, denoting the set of identifiers of spot and chained orders, respectively.[5]

The production assignment $\mathcal{P} : \mathcal{N} \to \mathcal{N}$ is a bijective function consisting of a single cycle. $\mathcal{P}(i)$ is the input in the set $\mathcal{N}$ that fulfills order i. Without loss of generality[6],

$$\mathcal{P}(i) = i + 1 \quad \text{for } i < N, \qquad \mathcal{P}(N) = 1.$$

[5] $\mathcal{N}^s$ and $\mathcal{N}^x$ are a partition of $\mathcal{N}$: $\mathcal{N}^s \cap \mathcal{N}^x = \emptyset$ and $\mathcal{N}^s \cup \mathcal{N}^x = \mathcal{N}$ with $N^s = |\mathcal{N}^s|$, $N^x = |\mathcal{N}^x|$.

[6] The numbering is a normalization under the single-cycle assumption.

An assignment specifies both the production ordering and the location of funded nodes: $\mathcal{D} = (\mathcal{P}, \mathcal{N}^s)$. With the production ordering normalized, organizational differences are represented by the placement of funded nodes along the circle. This organization is consequential, as will be clear.

Funding follows a simple rule. A funded order, $i \in \mathcal{N}^s$, has independent financing. Instead, a chained order $i \in \mathcal{N}^x$ is funded when good i is sold. That is, a chained order j can be executed only after its custom's order $\mathcal{P}^{-1}(j)$, has paid. Payments therefore propagate along the production circle. When order i pays for input $j = \mathcal{P}(i)$ and $j \in \mathcal{N}^x$, the seller of j uses the proceeds to finance order j, which pays for $\mathcal{P}(j)$, and so on, until the payment reaches a funded node. Funded nodes thus split the production network into a collection of payment chains, each starting from a single source, a node in $\mathcal{N}^s$. We assume $N^s > 0$, so every transaction is executed. Because all transactions are for the same amount, there is no problem of payments remaining stuck.

Payments-Chain Network. Given $\mathcal{D} \equiv (\mathcal{P}, \mathcal{N}^s)$, the payment-chain network is obtained by removing production links that enter funded nodes. Formally:

Definition 1. *A **payments-chain network** $\mathcal{K}$ is an acyclical directed network with nodes $\mathcal{N} = \{1, 2, ..., N\}$ and links $\mathcal{V} = \{(i, j) | \mathcal{P}(i) = i + 1 = j, j \in \mathcal{N}^x\}$, $\mathcal{K} = (\mathcal{N}, \mathcal{V})$. A **payments chain** of length n is a finite sequence of nodes $\{i_k\}_{k=0}^n$ such that the sequence starts with some $i_0 \in \mathcal{N}^s$ and $\forall k \in \{1, \ldots, n\}$, $\mathcal{P}(i_{k-1}) = i_k \in \mathcal{N}^x$. By convention, if $i \in \mathcal{N}^s$ and $\mathcal{P}(i) \in \mathcal{N}^s$, the sequence $\{i\}$ is a chain of length zero.*

The links in the payment chain network $\mathcal{K}$ record the onward flow of funds through chained nodes; funded nodes have no incoming link. Notice that the payment network simply removes links in the production network (associated with $\mathcal{P}$) directed to funded nodes. Thus, the payment chain network splits the single production cycle into a set of acyclic payment chains (sub-networks).

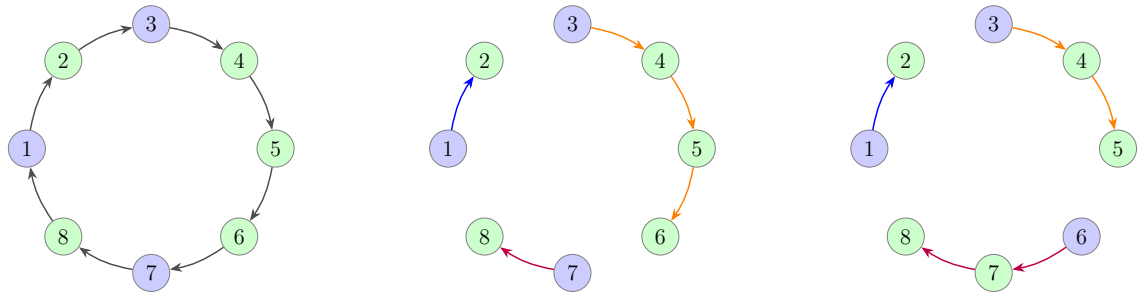
We define an order's **position**, n, in a payment chain as the number of payments that must be executed before an order is funded. Naturally, spot orders have position $n = 0$, chained orders receiving payment from spot orders have position $n = 1$, chained orders receiving payment from orders in position $n = 1$ have position $n = 2$, and so on.

The network admits a matrix representation. Let $\mathbf{s} \in \{0, 1\}^N$ indicate funded nodes, with $s_i = 1$ if $i \in \mathcal{N}^s$, and define the production matrix $\mathbf{P}$ with typical element by $P_{ji} = 1$ if $j = \mathcal{P}(i)$, and zero otherwise—ones in the subdiagonal and in the upper-right corner. The payment-chain network is represented by $\mathbf{V}(\mathbf{s}) = \big[\mathbf{I} - \text{diag}(\mathbf{s})\big]\mathbf{P}$. Notice

that $\mathbf{V}(\mathbf{s})^n\mathbf{s}$ identifies the orders at position n.[7] Hence, the fraction of orders occupying that position is $H(n;\mathcal{D}) = \mathbf{1}^\top \mathbf{V}(\mathbf{s})^n\mathbf{s}/N$.

Because each chain of length at least n contains exactly one order at position n. Thus, the difference between the numbers of orders at positions n and $n+1$ counts chains of length exactly n, i.e., $N\big[H(n;\mathcal{D}) - H(n+1;\mathcal{D})\big]$. Dividing by the total number of chains, N^s, gives the associated chain-length distribution, $g(n;\mathcal{D}) = N\big[H(n;\mathcal{D}) - H(n+1;\mathcal{D})\big]//N^s$.

Examples. Consider an example consisting of eight nodes. Let $\mathcal{N}^s = \{1,3,7\}$. Panel (a) of Figure 1 shows this production and funding assignment. Removing links $2 \to 3$, $6 \to 7$, and $8 \to 1$ yields the payment chains $1 \to 2$, $3 \to 4 \to 5 \to 6$, and $7 \to 8$, with lengths $(1,3,1)$, as in panel (b). In this example, order 6 occupies position $n = 3$, while orders 2, 4, and 8 occupy position $n = 1$, etc. Panel (c) illustrates a different assignment: keeping the production ordering unchanged, funded nodes at $\{1,3,6\}$ yield chains of lengths $(1,2,2)$, yielding a different distribution of lengths and positions.



(a) Production relation **(b) Random assignment** **(c) Efficient assignment**

Figure 1: Production and Funding Assignments

Notes: Blue nodes are funded; green nodes are chained. Panel (a) shows the production network. Panels (b) and (c) show two payment-chain networks, corresponding to a random and an efficient assignment. Colored arrows distinguish payment chains.

TFP. As documented in our empirical analysis below, not only do we observe a number of transactions whose settlement is contingent on prior transactions settling, but we also observe transaction delays, and non-trivial economic losses for those who don't pay on the spot. Thus, the payments-chain network must determine who and by how much bears those losses.[8]

[7] The vector of node positions is $\mathbf{n} = \sum_{k=1}^{N-1} \mathbf{V}(\mathbf{s})^k\mathbf{1}$, where $\mathbf{1}$ is the unit vector and n_i is the position of node i in its corresponding chain.

[8] The timing-based estimates preserved in Appendix C.1.3 imply delays of between 5-25 business days a year for the last firm in the average chain, equivalent to a delay and productivity loss of between 2% and 10% per year, with longer chains being associated with longer delays.

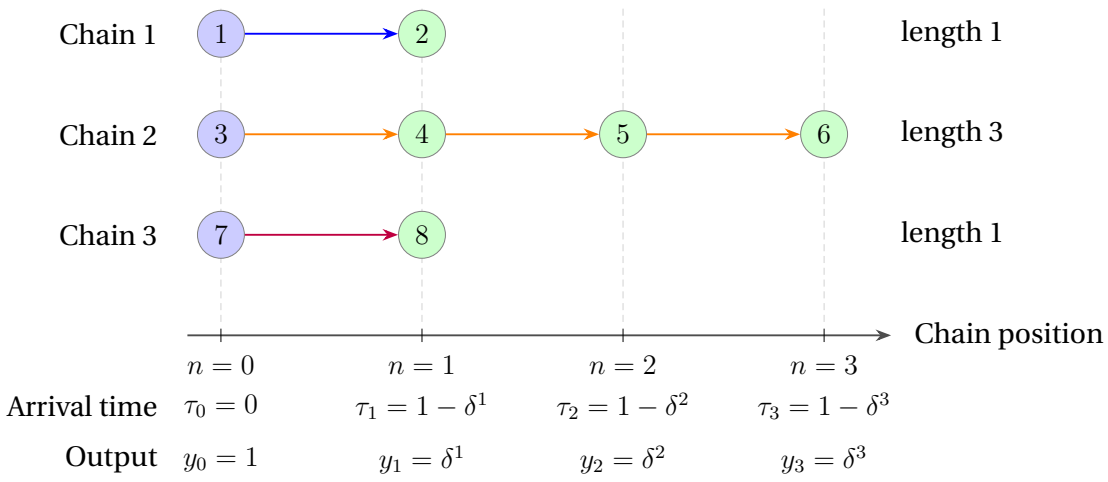


Figure 2: Payment Chains, Arrival Times, and Output

Notes: The chains from Figure 1(b) are aligned by position. Horizontal spacing represents chain position, arrival times, and output following the capacity-utilization mechanism.

From the perspective of our theory, the payment network would be inconsequential if the transfer of funds across the entire network were instantaneous. In the extreme, a single dollar, i.e., a single spot order, could fund the entire economy's transactions if it could be passed on and settled immediately.

To capture the timing of payments, let $\tau(n)$ denote the time at which payment is executed for position n, with $0 = \tau(0) < \tau(1) < \tau(2) < \cdots < 1$. Transaction delays can arise due to limited technology to process transactions, e.g., the reliance on physical currency transfers or the clearing time of checks and electronic payments. They may also serve as a withholding protocol to prevent fraud. Now, even if transfers are not instantaneous, the network would be inconsequential if no resources were lost due to payment speed. To capture the economic losses from delays, let y_n denote output per transaction relative to full utilization, equivalently capacity utilization, where $y(0) = 1 \geq y(1) \geq y(2) \geq \cdots > 0$.[9]

We specify the settlement times τ and output function, $y(n)$, following from a transaction technology and capacity-utilization mechanism. Assume that each order corresponds to intermediate inputs, divisible into a continuum of units, which the buyer can process at a constant rate over the period $[0, 1]$. An order at position n receives its input at time τ_n. Before delivery, the buyer's processing capacity remains idle (i.e., its labor or capital is underutilized). After the delivery of the input, the buyer processes inputs during the remaining fraction of the period. Hence, output and capacity utilization is

[9]Transaction i, at position n_i, contributes y_{n_i}/N to normalized aggregate output. Summing across transactions gives $N^{-1}\sum_{i=1}^{N} y_{n_i} = \sum_{n=0}^{\infty} H(n;\mathcal{D})y_n$. Note that if $y_n = 1$, final output equal the aggregate value of inputs, as with a linear technology.

$$y_n = 1 - \tau_n.$$

Payment delays therefore reduce output by shortening the time over which existing productive capacity can be utilized.

We microfound the delays as follows. Before releasing payment to the supplier, the buyer must process and inspect a fraction $1-\delta$ of the input units it will actually use, as a form of quality control to deter fraud—inspected units are still turned into output. Since the buyer uses y_n units, verification takes $(1-\delta)y_n$ units of time. Once verification is complete, payment is released and becomes available to finance the next order in the chain:

$$\tau_{n+1} = \tau_n + (1-\delta)y_n = \tau_n + (1-\delta)(1-\tau_n), \qquad \Rightarrow \qquad y_{n+1} = 1 - \tau_{n+1} = \delta\, y_n.$$

With $\tau_0 = 0$ and $y_0 = 1$, it follows that $y_n = \delta^n$. This geometric output function thus arises from sequential verification and utilization losses. Appendix A.1 provides a fuller description based on protection against two-sided fraud.

The geometric output function admits other interpretations. It can simply correspond to a period discounting in a setting where each transaction takes one period, $\tau(n) = n$. In that case, δ^n is the time-zero present value. This is exactly the interpretation in the celebrated Townsend Turnpike model (see Townsend, 1980) or in the first formulation of money-search models (see Kiyotaki and Wright, 1989).

Figure 2 organizes the nodes corresponding to the payment chains in terms of their positions. The arrival times and output shown under each indicate the arrival of funds and the corresponding production.

For the results on aggregate production that follow, the geometric output function is inessential. If each transaction takes a fixed amount of time, $1/\eta$, then $\tau_n = n/\eta$ and utilization time runs out linearly, $y_n = 1 - n/\eta$; we obtain similar qualitative outcomes.[10] We adopt $y_n = \delta^n$ for two reasons: it delivers more elegant formulas and allows us to settle all transactions within a unit-time interval in between the expenditure decisions in the infinite-horizon business-cycle setting.

Chain-Length Distributions and Aggregate Output. Let $\mathbf{y}$ collect output across orders. Funded orders produce one unit, while each chained order produces δ times the output associated with its predecessor. Thus, $\mathbf{y} = \mathbf{s} + \delta \mathbf{V}(\mathbf{s})\mathbf{y}$, yielding[11]

[10] Orders at or beyond position η cannot be served within the period. Appendix B.4 studies a related truncation with geometric output up to a cutoff.

[11] We can replace finite sums with infinite sums because the payment network is acyclic.

$$\mathbf{y} = \sum_{n=0}^{\infty} \delta^n \mathbf{V}(\mathbf{s})^n \mathbf{s} = \left[\mathbf{I} - \delta \mathbf{V}(\mathbf{s})\right]^{-1} \mathbf{s}.$$

Define the share of unfunded orders, $\mu = N^x/N$. Aggregate output per order and average output among chained orders are therefore

$$\mathcal{Y}(\mu, \mathcal{D}) = \frac{\mathbf{1}^\top \mathbf{y}}{N}, \qquad \mathcal{A}(\mu, \mathcal{D}) = \frac{(\mathbf{1} - \mathbf{s})^\top \mathbf{y}}{N^x},$$

Since funded orders produce one unit, aggregate output satisfies $\mathcal{Y}(\mu, \mathcal{D}) = (1 - \mu) + \mu\, \mathcal{A}(\mu, \mathcal{D})$.

Differences in output reflect differences in capacity utilization and thus measured TFP.

A novel aspect of our theory is that the number and placement of funded nodes therefore affects TFP. As a baseline, we compare two funding assignments, leaving other variations to the extensions below. Under a **random assignment**, the N^s funded nodes are taken uniformly from the set of permutations of N^s in the N. Under the **efficient assignment**, the vector $\mathbf{s}$ is chosen to maximize output $\mathbf{1}^\top \left[\mathbf{I} - \delta \mathbf{V}(\mathbf{s})\right]^{-1} \mathbf{s}$ subject to $\mathbf{1}^\top \mathbf{s} = N^s$.[12] Next, we obtain the asymptotic efficiency for the random and efficient assignments.

Asymptotics. We now derive moments of the payment-chain network as $N \to \infty$, while keeping ratio $\mu \equiv N^x/N$ constant. For the two assignments and the geometric output function derived above, the distribution of chain lengths and TFP aggregate is as follows:

Proposition 1 (Random versus Efficient Networks)**.** *Given a fraction μ of chained orders and output function $y_n = \delta^n$, as $N \to \infty$, we obtain the following solutions:*

Assignment	**Auxiliary Objects**	**Distribution** $g(n; \mu, \mathcal{D})$	**Chained Output** $\mathcal{A}(\mu; \delta, \mathcal{D})$
Random	–	$(1-\mu)\mu^n$	$\delta \frac{1-\mu}{1-\delta\mu}$
Efficient	$\underline{n} = \lfloor \mu/(1-\mu) \rfloor$ $\overline{n} = \underline{n} + 1$ $p = \overline{n} - \mu/(1-\mu)$	$\begin{cases} p \text{ if } n = \underline{n} \\ 1-p \text{ if } n = \overline{n} \\ 0 \text{ otherwise} \end{cases}$	$\delta \frac{1-[p\delta^{\underline{n}} + (1-p)\delta^{\overline{n}}]}{1-\delta} \frac{1-\mu}{\mu}$

$\mathcal{A}$ and $\mathcal{Y}$ are weakly (strictly) decreasing in μ for the random (efficient) assignment and increasing in δ . Moreover,

$$\lim_{\mu \to 0} \mathcal{A}(\mu; \delta) = \delta, \quad \lim_{\mu \to 1} \mathcal{A}(\mu; \delta) = 0, \quad \lim_{\delta \to 0} \mathcal{A}(\mu; \delta) = 0, \quad \lim_{\delta \to 1} \mathcal{A}(\mu; \delta) = 1.$$

[12] In the literature, the transactions in Kiyotaki and Wright (1989) correspond to a random assignment whereas Townsend (1980) corresponds to an efficient assignment.

Output per order, $\mathcal{Y}(\mu) = (1-\mu) + \mu\mathcal{A}(\mu)$, is concave in μ.

In the random assignment, the geometric distribution emerges from random funding placement. Starting at a funded node, a chain has length n if the next n nodes are chained and the following node is funded. In the limit $N \to \infty$, these events have probability $\mu^n(1-\mu)$, so $g(n;\mu) = (1-\mu)\mu^n$. The distribution of positions is also geometric, $H(n;\mu) = (1-\mu)\mu^n$. The formula for output follows immediately from $y_n = \delta^n$.[13]

The efficient assignment minimizes output losses by equalizing chain lengths as much as the integer nature of the chains permits. To see why this is optimal, consider any assignment with chains of varying lengths. If we have chains of lengths n_1 and n_2 such that $n_1 + 1 < n_2$, we could move an order from position n_2 in the longer chain to extend the shorter chain to length $n_1 + 1$. This reallocation strictly increases total output since the relocated order gains $y_{n_1+1} - y_{n_2} = \delta^{n_1+1} - \delta^{n_2} > 0$. Panels (b) and (c) in Figure 1 provide a concrete example: moving the final node in the chain of length 3 to the final position in a chain of length 1 allows that node to increase its output by $\delta^2 - \delta^3 > 0$ without causing losses in any other node. In general, such improvements are possible whenever chains differ by more than one in length. Therefore, the efficient assignment must concentrate all chains on at most two consecutive integers: chains of length $\underline{n}$ and $\overline{n} = \underline{n}+1$. The weight p on chains of length $\underline{n}$ ensures that the average ratio of chained to spot orders across chains yields the ratio across all orders $\mu/(1-\mu)$.

Under both assignments, $\mathcal{A}$ is decreasing in μ with well-behaved limits, and aggregate output $\mathcal{Y}$ is concave in μ.[14] These properties will reemerge in the business-cycle setting.

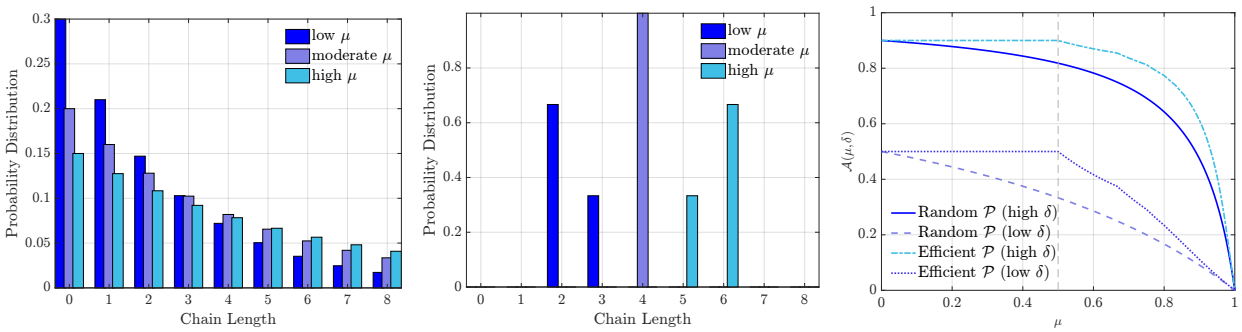
Figure 3 illustrates the key differences between random and efficient assignments: Panel (a) shows the geometric distribution of chain lengths following random assignment for three values of μ. Greater μ induces longer chains. Panel (b) shows the efficient assignment counterpart exhibiting distributions concentrated on only two adjacent lengths. Panel (c) compares $\mathcal{A}(\mu)$ across both assignments for $\delta \in \{0.5, 0.9\}$. Clearly, the efficient assignment outperforms the random assignment.[15]

This section's main insight is that economic losses can result from payment de-

[13] Note $\mathcal{Y} = \sum_{n=0}^{\infty} H(n;\mu)\cdot y_n = \sum_{n=0}^{\infty}(1-\mu)\mu^n \cdot \delta^n = \frac{1-\mu}{1-\mu\delta}$.

[14] As $\mu \to 1$, chains become arbitrarily long and the average final goods per chained order tends to zero. Conversely, as $\mu \to 0$, a chained order is guaranteed to end in position 1, obtaining $\mathcal{A} = \delta$ in expected goods. Similarly, as $\delta \to 0$, chained orders receive nothing, whereas any losses vanish as $\delta \to 1$. Under the efficient assignment, $\mathcal{Y}$ is piecewise linear in μ, with kinks where $\mu/(1-\mu)$ is an integer and slopes that fall at each kink. Between kinks, $\mathcal{A}$ itself is convex in μ, so concavity is a property of $\mathcal{Y}$ rather than of $\mathcal{A}$; see Appendix A.2.2.

[15] Under the efficient assignment for $\mu \leq 1/2$, $A(\mu) = \delta$ because it is possible to assign each chained order to a spot order.



(a) g(n) (random) **(b) g(n) (efficient)** **(c) Output $\mathcal{A}(\mu)$**

Figure 3: Chain-Length Distributions and Average Output per Chained Order

Notes: Panels (a) and (b) show distributions $g(n;\mu)$ for $\mu \in \{0.7, 0.8, 0.85\}$ under random and efficient assignments. Panel (c) shows $\mathcal{A}(\mu)$ for both assignments with $\delta \in \{0.5, 0.9\}$.

lays and the possible disorder of payments. It is worth distinguishing the sources of these losses from other environments. Unlike search models, there are no congestion externalities—the production assignment is one-to-one. Unlike sticky-price models, prices have no allocative implications as long as goods are transferred. Instead, economic losses arise exclusively from delays, which is clear as the economy is at full capacity when all transactions are spot or $\delta \to 1$. This is true under both payment assignments, but importantly, losses are magnified when payments are disorganized. The efficient assignment minimizes these losses by organizing all chains at similar lengths.

2.1 Interpretation and Extensions

The random and efficient assignments are polar cases. The random assignment describes a decentralized market in which a Walrasian mechanism clears the market for inputs but does not organize who pays whom, so agents placing chained orders are uncertain about when funds arrive and goods are delivered. The efficient assignment describes an economy in which a clearinghouse or similar entity sequences payments to minimize delays—an arrangement that requires substantial coordination. Actual payment networks likely lie in between. Appendix B studies several variations of the environment: correlated assignments, in which spot orders are more or less likely to be linked to other spot orders (unit entries in $\mathbf{s}$ are positively correlated); a pooled assignment, in which agents placing many chained orders pool incoming funds so that none remain idle; costly reassignment of orders in long chains; default cascades, in which orders beyond a threshold position are abandoned; partial delays, in which some transactions settle instantly; and alternative uses of capacity while deliveries are pending.

Each extension changes the mapping $\mathcal{A}(\mu)$, but not the essence.

Discussion: Reduction of Economic Complexity. We clearly abstract from more realistic features for tractability, enabling the isolation of core mechanisms. Real economies exhibit greater complexity: transactions vary by size (100 vs 200), production networks involve multilateral relationships (the production network is more complex), and payment timing is more nuanced.[16] The random assignment captures this inherent complexity, suggesting that asymptotic approach may be more realistic than the assumption of perfect optimization. Indeed, the pooled and efficient assignments represent increasing levels of coordination: from simple income pooling to complete payment-network optimization. The goal so far is to provide a simple framework to speak to these issues. Section 4 leverages the aggregation to analyze agents' endogenous decisions between placing spot and chained orders.

Another feature is the constant price assumption. That is, we abstract from strategic pricing decisions. In fact, there's no obvious direction: should prices result from sequential bargaining in chains or cooperative games a la Shapley? We think that our formulation is a good starting point: as long as the timing of transactions is what matters for production, prices would affect the transfer of funds across parties but would not affect productive efficiency.

2.2 From Payments-Chain Networks to Portable Budget Constraints

We now demonstrate how the transactions and delays in the payment chain aggregate into reduced-form budget constraints when agents place many orders. We can embed this budget representation into standard models.

Timing. The goal is to arrive at a budget representation allowing agents to make expenditure and financial decisions at discrete dates $t = 0, 1, \ldots$, with transactions clearing within the unit interval between t and $t+1$. That is, the transaction times $\tau(n)$ encountered above happen between t and $t+1$. More explicitly, we could write $\tau_t(n) = t + \tau(n)$ as the calendar time of a transaction agreed at t that lands in position n but we will not need this level of detail.

We note that this timing is akin to the alternation between Walrasian markets and bilateral trade in Lagos and Wright (2005), with the distinction that, here, each agent

[16] When transactions differ in size, an agent receiving \$100 but owing \$200 faces complex cash flow problems that require sophisticated optimization—challenges that even large firms struggle with, analogous to NP-hard problems in logistics. We anticipate that methods from statistical physics could be used to address the combinatorial challenges arising from transactions of different sizes.

undertakes many interconnected transactions, which follow the sequence dictated by the network. The key is that transactions settle by the next decision date, since $\tau(n) \in [0,1]$, avoiding the need to carry pending orders across periods. Taking N to be large, eliminates position risk and allows aggregation.

Individual Transactions. Recall that under price normalization above, each transaction has value $p = 1/N$ for N goods in the overall economy, normalizing the value of inputs to 1. Consider agent a sells the subset of inputs $\mathcal{S}_a \subseteq \mathcal{N}$ and places orders for the subset of inputs $\mathcal{O}_a \subseteq \mathcal{N}$. Let $\mathcal{O}_a^s = \mathcal{O}_a \cap \mathcal{N}^s$ and $\mathcal{O}_a^x = \mathcal{O}_a \cap \mathcal{N}^x$, with counts $M_a^y = |\mathcal{S}_a|$, $M_a^s = |\mathcal{O}_a^s|$, and $M_a^x = |\mathcal{O}_a^x|$. Its sales receipts and expenditures are

$$y_a = \frac{M_a^y}{N}, \qquad E_a^s = \frac{M_a^s}{N}, \qquad E_a^x = \frac{M_a^x}{N}.$$

Each chained order is financed by the sale of its corresponding input, so $\mathcal{O}_a^x \subseteq \mathcal{S}_a$, and those receipts remain committed to that order. We suppress the agent index on scalar variables below.

Within-Period Budgets. All transactions settle within each discrete period t. Let n_i denote order i's chain position. Input j is purchased by order $\mathcal{P}^{-1}(j)$, so its sale occurs at position $u_j = n_{\mathcal{P}^{-1}(j)}$; a chained purchase i occurs at position $o_i = n_i$. Let W_n denote financial wealth after payments by orders at position n, starting from wealth W. Then

$$\begin{aligned} W_0 &= W + \frac{1}{N}\sum_{j\in\mathcal{S}_a} \mathbb{I}[u_j = 0] - E^s, \\ W_n &= W_{n-1} + \frac{1}{N}\sum_{j\in\mathcal{S}_a} \mathbb{I}[u_j = n] - \frac{1}{N}\sum_{i\in\mathcal{O}_a^x} \mathbb{I}[o_i = n], \qquad n \geq 1. \end{aligned}$$

Each input is sold and each order is settled once. Summing therefore gives end-of-period wealth $W^{\text{end}} = W + y - E^s - E^x$. Writing $W^{\text{end}} = W'/R$, where R is the gross return between periods, yields

$$\frac{W'}{R} + E^s + E^x = y + W. \tag{1}$$

Diversification and the Period Budget. With finitely many chained orders, unknown positions create idiosyncratic risk in final goods receipts, complicating dynamic decisions. We therefore consider many small transactions: as $N \to \infty$, hold the agent's endowment and expenditures fixed while transactions become smaller and more nu-

merous. Assume that the empirical distribution of its chained positions converges in probability to the network's conditional distribution of chained positions. For $E^x > 0$, final goods obtained from chained orders satisfy

$$X = \frac{1}{N}\sum_{i\in\mathcal{O}_a^x} \delta^{o_i} = E^x \left(\frac{1}{M^x}\sum_{i\in\mathcal{O}_a^x} \delta^{o_i} \right) \xrightarrow[N\to\infty]{p} \mathcal{A}(\mu,\mathcal{D})E^x.$$

Each order yields δ^{o_i}/N units of final goods because its input lot has size $1/N$, and the second equality uses $1/N = E^x/M^x$. Recall that, holding E^x fixed, smaller input lots imply more chained orders, $|\mathcal{O}_a^x| = M^x = NE^x$, leaving their total input quantity $M^x/N = E^x$ unchanged.[17] The convergence reflects the "big family" diversification assumption in the tradition of Lucas and Stokey (1987) and Shi (1997).

For a given assignment $\mathcal{D}$, define $q(\mu) = \mathcal{A}(\mu,\mathcal{D})^{-1}$. Since spot orders yield $S = E^s$, the limiting budget is

$$\frac{W'}{R} + S + q(\mu)X = y + W. \tag{2}$$

The effective price q captures the output loss from payment delays. Diversification preserves the order-specific funding commitments; redirecting receipts across pending orders instead corresponds to the pooled assignment in Appendix B.2, with its associated productivity factor and effective price. We can incorporate this budget constraints into various settings.

Pecuniary Externality. The aggregation of budget constraints also reveals that, in addition to possible inefficiencies from the assignment of orders, the environment features a pecuniary externality when agents take $q(\mu)$ as given. Indeed, $q(\mu)$ is the average, not a marginal cost of buying goods through chained orders. This cost depends on μ, affected by individual decisions but not internalized.

As anticipated, one way to eliminate the externality is to centralize transactions through an intermediary that aggregates all orders, treating the timing of transactions, $\tau(n)$, and the output function, $y(n)$, as if these were part of a production technology. That intermediary would treat orders—and their corresponding inputs—as the inputs $\{E^s, E^x\}$ of an aggregate production function:

$$\mathcal{F}(E^s,E^x) \equiv \mathcal{Y}\times(E^s+E^x) = E^s + E^x\mathcal{A}\left(\frac{E^x}{E^s+E^x},\mathcal{D}\right).$$

In that case, the intermediary would re-establish efficient pricing by setting $\mathcal{F}_{E^s}$ and

[17] Think of wheat: the same quantity measure of wheat can be divided into bushels, kilos, or grains; smaller units mean more units, with the total quantity unchanged.

$\mathcal{F}_{E^x}$ as the price of spot and chained orders, respectively. Of course, how realistic it is to assume centralization of payments to establish efficiency is in the eye of the beholder. Thus, in the business-cycle setting below, we contrast the economy with the pecuniary externality against a planner's solution that avoids it.

Taking Stock. The analysis so far provides a tractable mapping from payment chains to settlement delays, effective prices, and aggregate productivity. Before embedding this block into a dynamic business-cycle environment, the next section presents the motivating evidence.

3 Evidence

We now turn to the data to examine whether payment chains are an important feature of interfirm trade and whether disruptions to these chains affect real activity. Guided by the framework in Section 2, we focus on orders that are not paid upfront and ask whether their fulfillment waits for settlement. The key distinction from conventional trade credit is whether deferred payment finances goods already delivered or instead leaves orders unfulfilled until payment. We also examine whether firms' payments to suppliers systematically follow receipts from their own customers. To identify the real consequences of payment disruptions, we exploit the unexpected freezing of customer accounts at Financiera Desyfin. We show that sellers awaiting payment from affected buyers suffer larger declines in TFP, sales, and value added.

Obtaining this evidence poses substantial measurement and identification challenges: While some recent datasets contain firm financials and invoices, they typically lack information on payment dates and product-level delivery timing. Costa Rica offers a particularly rich setting: its administrative data allow us to link the universe of electronic invoices and interbank business payments to measures of delivery timing and firm outcomes. This richness, combined with the natural experiment, allows us to provide evidence of chained payments, a previously unexplored dimension of interfirm trade, and to identify the real consequences of their disruption.

Data. We combine two administrative data sources. The first is Sinpe—Costa Rica's National Electronic Payments System—which processes the universe of interbank payments.[18] In 2024 alone, Sinpe recorded over 62 million transactions. For each transaction, we observe the timestamp, sender and receiver identifiers, banks involved, and

[18]Sinpe is the acronym for Costa Rica's National Electronic Payments System (*Sistema Nacional de Pagos Electrónicos*), operated by the Central Bank of Costa Rica.

transferred amounts. The coverage spans interbank activity for 243,234 distinct firms, with an average payment of USD 4,507.[19]

The second is Costa Rica's VAT reporting system which encompasses the electronic invoices for every formal sale. Each invoice records the sale price, product code, quantity sold, and flags whether payment was settled on the spot or was deferred. Deferred status is recorded when the legally valid electronic invoice is issued, together with the corresponding contractual repayment window. Sellers have incentives to correctly classify a payment as deferred, as this allows the seller to postpone the VAT remittance of the transaction's unpaid balance for up to 90 days. Because these terms are recorded at the contracting date, they offer time-stamped information on repayment terms and the deadline for full repayment at a level of granularity rarely available.

The invoice and bank payment records play complementary roles. Invoices with deferred payment identify transactions with the first feature of chained orders, i.e., that they are not paid upfront. Moreover, we link invoices' payment deadlines to payments for the associated delivery and transportation services. Combining payments records with logistics-payment dates and the timing of downstream product sales allows us to examine whether fulfillment is linked to payment timing. The economic distinction is whether deferred payment finances orders awaiting fulfillment or goods already delivered, which distinguishes chained orders from conventional trade credit.[20]

The combination of invoice and bank payment data allow us to further identify whether the buyer's realized payment behavior is contingent on receiving payments, the third feature of chained orders. In particular, we identify buyers whose payments to suppliers systematically follow payments received from their own customers.

We proceed by providing evidence that chained orders are used systematically and have economic consequences. Fact 1 shows that deferred payment is widespread and is used disproportionately by buyers whose outgoing payments closely follow recent receipts. Fact 2 shows that deferred-payment orders often remain unfulfilled *until the buyer is due to pay in full*, supporting the timing mechanism at the heart of the model. Fact 3 documents a substantial price premium associated with deferred terms within narrowly defined product and trading relationships. Finally, Fact 4 uses an unexpected freeze of bank accounts to show that interrupting a scheduled payment reduces sellers' real activity, with the largest effects when the payment was due soon after the freeze.

[19]For more details on Costa Rica's firm network, see Méndez and Van Patten (2025); Argente et al. (2025). Alvarez et al. (2023) focuses on Sinpe Móvil, a peer-to-peer mobile money system.

[20]Appendix C.1 describes the payment-timing measure and reports extensive- and intensive-margin estimates, amount-based tests, dynamic responses, and future-receipt placebos.

3.1 Fact 1: Receipt-dependent payments are widespread

Our data show that deferred payments are common in firm-to-firm trade accounting for approximately 60 percent of the universe of transactions. Table 1 reports the distribution of repayment windows. Conditional on being a deferred-payment invoice, the average repayment window is 25 days (median is 30 days) and 10% of invoices have a repayment window of a month and a half. Figure 4 shows that deferred payment is not concentrated within one industry, but widespread, and accounts for a substantial share of transactions in certain sectors. While deferred payments may be observed when the supplier is providing trade credit, below we show that for a large share of the deferred-payment transactions, the products only arrive once full payment is due.

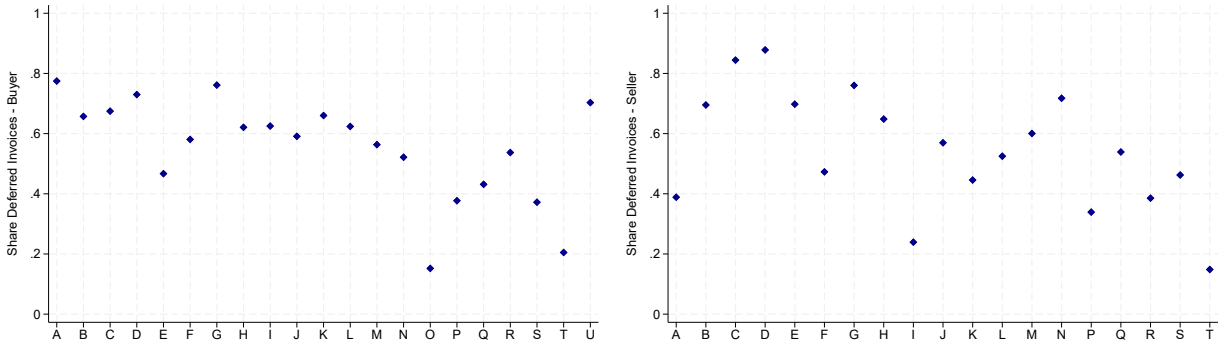
Table 1: Distribution of Repayment Windows

Repayment window in days

Statistic	10th pct.	25th pct.	Median	75th pct.	90th pct.	Mean
Days	7	8	30	30	45	25.13

Notes: The table reports the distribution of the repayment window for deferred-payment invoices. The repayment window is measured in days.

Figure 4: Share of Deferred Transactions by Buyer's and Seller's Industry



(a) Buyer's industry **(b) Seller's industry**

A: Agric., forestry & fishing	**H:** Transport & storage	**O:** Public admin. & defense; social security
B: Mining & quarrying	**I:** Accommodation & food services	**P:** Education
C: Manufacturing	**J:** Information & communication	**Q:** Health & social work
D: Electricity, gas, steam & AC	**K:** Finance & insurance	**R:** Arts, entertainment & recreation
E: Water, sewerage, waste & remediation	**L:** Real estate	**S:** Other services
F: Construction	**M:** Professional, scientific & technical	**T:** Household employers & own-use production
G: Wholesale/retail; motor-vehicle repair	**N:** Admin. & support services	**U:** Extraterritorial organizations

Notes: The figures show the share of transactions with a deferred payment by industry. Panel (a) classifies transactions based on the buyer's industry, while Panel (b) does so based on the seller's industrial classification.

Incoming payments predict subsequent outgoing payments. A defining feature of a payment chain is that a firm's payment to a supplier depends on the arrival of payment from its own customer. We ask whether this behavior is visible directly in the bank-transfer data, without using invoice terms. For each buyer, we measure the increase in the probability of observing a payment to a supplier after observing an incoming payment from a customer. Calendar-date fixed effects absorb aggregate payment cycles, including month-end settlement, tax dates, holidays, and banking disruptions. Firm-by-weekday fixed effects absorb each firm's regular payment schedule. Thus, the coefficient is identified from deviations from the firm's usual behavior on the same weekday.

An incoming payment raises the probability of an outgoing payment during *the next business day* by 3 percentage points, or 7% relative to the dependent-variable mean (column (1) of Table C.1). In line with the model's intuition, the response is larger among small firms and attenuates as asset-to-sales or liabilities-to-sales ratios rise (columns (2)-(4) of Table C.1). These outgoing payments are *not* triggered if future receipts are used as placebos (Table C.4). Moreover, in line with the payments-chain mechanism, the effect is driven by instances in which the incoming payment is large enough to fund the outgoing payment (Table C.2). These patterns are hard to explain through a regular payment calendar or persistent differences in scale.

We use this payment-timing pattern to construct a firm-level classification. In each month, a firm is classified as chained when its payments to suppliers are especially likely to follow recent payments received from customers; otherwise, it is classified as spot. Appendix C.1.3 describes the classification in detail. Under this definition, 45% of firms are classified as chained. Buyers classified as chained are 26% more likely to use deferred terms—a 7 percentage point difference—and their deferred invoices have longer repayment windows: approximately 34 days on average, compared with 26 days among other buyers. Because the classification is constructed from bank transfers, whereas deferred status and repayment windows come from electronic invoices, this correspondence is not mechanical. Appendix C.1 reports the underlying payment-timing regressions and robustness exercises.

Taking Stock. The two independently constructed measures point to the same underlying payment behavior. Electronic invoices record whether payment is deferred and when full payment is due, while bank transfers identify buyers whose payments to suppliers tend to follow receipts from customers. *We use this measure of incoming–outgoing payment dependence to test whether Facts 2–4 also emerge under an independent, timing-based definition of chained behavior.* Deferred-payment invoices remain

our baseline measure, and each fact below points to the corresponding evidence based on bank-payment timing.

3.2 Fact 2. Good deliveries often await payment

The economic meaning of deferred payment depends on what happens to the goods while payment is pending. An outstanding invoice can finance goods already in use or an order whose fulfillment is still waiting. Distinguishing these arrangements requires connecting payment terms to the subsequent movement and use of goods. Our linked records allow us to examine this margin through two complementary exercises. Across sectors, we study payments for delivery and transportation services around contractual invoice due dates. In retail, we follow newly purchased products into the buyer's own sales. The first exercise provides broad evidence from logistics payments; the second follows the purchased product into downstream economic activity. Together, they provide evidence on a distinction that payment terms alone cannot reveal and that is rarely observable.

To our knowledge, this is the first economy-wide attempt to connect contractual payment deadlines to a validated measure of delivery timing. Because direct physical delivery dates are not recorded at this scale, we construct a delivery-activity proxy from payments to delivery and transportation providers and validate its timing independently using the onset of retail sales of newly purchased products. Namely, for every deferred-payment invoice, we observe the date at which the repayment window closes. Separately, the electronic invoice data span the universe of firm-to-firm transactions, and therefore also identify the events in which one firm pays another for delivery or transportation services. We therefore estimate an event study across sectors that asks if the probability of observing a delivery-service payment rises after a deferred-payment invoice reaches its due date. We implement this exercise for deferred invoices issued to buyers that are observed purchasing delivery or transportation services during the sample period.[21] Formally, for deferred invoice v with due date d_v, we build event time $e = t - d_v$ and estimate

$$DeliveryPayment_{vt} = c + \sum_{e \neq -1} \theta_e \mathbf{1}\{t - d_v = e\} + \alpha_{b(v)} + \lambda_{p(v)} + \gamma_t + u_{vt}, \tag{3}$$

where $DeliveryPayment_{vt}$ indicates that the buyer in transaction v makes a payment to a delivery or transportation provider in week t, $b(v)$ and $p(v)$ denote the buyer and

[21] Thus, firms that never pay for transport do not enter the estimation as permanent zeros; the coefficients are identified from the timing of logistics payments among firms for which such payments are observable.

13-digit product, respectively, and γ_t denotes calendar-week fixed effects. We omit the week immediately preceding the due date, so $\theta_{-1} = 0$. The exercise is designed to measure buyer-level logistics activity around each invoice's contractual due date.[22]

Figure 5 plots the event-time profile in levels.[23] This exercise is hard to implement in almost any other setting because it requires observing contractual repayment windows and economy-wide delivery-service payments in the same administrative system. A rise at $e = 0$ directly connects the end of the repayment window with logistics activity, the timing margin at the heart of the model.

As shown in Figure 5, the probability of paying for delivery or transportation services rises from 5.3% in the week immediately preceding the due date to 23.7% in the due-date week, an increase of 18.4 percentage points. The figure also shows a rise roughly four weeks before the due date. Since the average repayment window is 25 days, this earlier increase is consistent with some deferred-payment transactions being delivered near the invoice date and paid later, as in conventional trade credit. The spike at the due date, however, shows that this is not the only pattern. For a substantial set of transactions, logistics activity is concentrated once full payment becomes due. This is the timing margin emphasized by the model: deferred settlement is not merely an accounting maturity date, but is often linked to when goods move through the supply chain.

Following purchased goods into downstream sales. Retail transactions provide a second perspective on fulfillment from the buyer's side, complementing the logistics evidence *by tracing the purchased product itself.* We ask whether buyers begin selling products purchased under deferred-payment terms when the invoice is issued or only after the payment becomes due. This requires matching products purchased from a supplier to those subsequently sold by the buyer. We therefore focus on retail firms, which resell goods without in-house transformation. A shoe retailer, for example, resells the same shoes it purchases from its suppliers. An automaker instead combines tires and other components into cars, so purchased inputs cannot be matched one-to-one to the products it sells.

Within retailers, we focus on newly purchased products with deferred payment and compare the onset of resale with invoice issuance and the buyer's payment deadline. In about 70% of these cases, *the product first appears in the retailer's sales after the retailer's*

[22] Namely, the exercise is not meant to require a one-to-one assignment of a particular carrier bill to a particular shipment. The retail exercise below supplies an independent product-level validation by asking whether this logistics activity coincides with the point at which a newly purchased product enters downstream sales.

[23] Hence, the point at $e = -1$ is the estimated baseline level rather than zero.

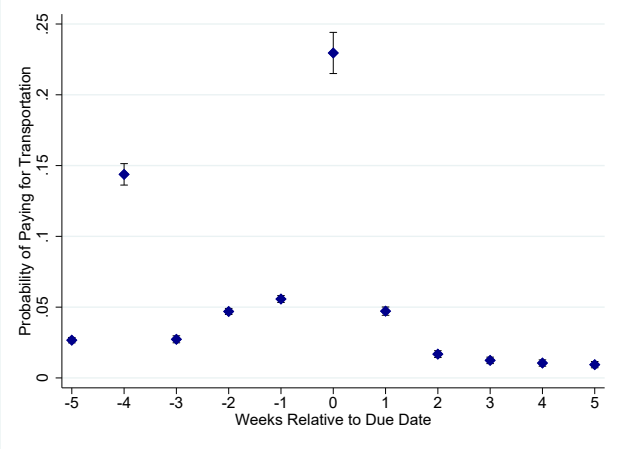


Figure 5: Delivery-Service Payments Around Deferred-Payment Due Dates

Notes: The figure plots the estimated event-time profile $100(\widehat{c} + \widehat{\theta}_e)$ from the linear probability model in equation (3). The week immediately preceding the due date is the omitted event-time category, with $\widehat{\theta}_{-1} = 0$; its plotted value is therefore the estimated baseline level. The regression includes buyer, 13-digit product, and calendar-week fixed effects. Standard errors are clustered at the buyer level, and error bars denote 95% confidence intervals. For computational tractability, the estimation uses a 20% random sample of invoices whose observation windows lie entirely within 2024.

payment to its supplier is due. Figure C.5 traces resale activity for such transactions around the payment deadline.

Validation for the delivery-timing measure. The onset of retail sales provides a powerful validation of shipping-payment dates. Our linked payment records and product-level invoices allow us to follow *new* products into retailers' own sales and compare the week in which a product first appears in downstream sales *with the timing of the associated shipping payment.* Importantly, we retain all newly purchased products, regardless of whether resale begins before or after that payment. Figure C.6 shows a sharp correspondence: in the exact week a new product first appears in the retailer's sales, the probability that the retailer makes a payment for delivery or transportation jumps by an amount equal to nearly 200% of its sample mean. No increase is larger at any horizon in the six weeks before or after product entry. This exact-week spike provides strong independent validation that shipping-payment dates track the timing at which goods enter downstream activity and become available for resale. Because direct delivery dates are generally unavailable at economy-wide scale, this correspondence provides rare direct evidence on delivery timing. We use this validated proxy to identify the transaction-level timing pattern most closely associated with chained fulfillment: deferred-payment transactions for which the associated delivery-service payment is not observed until the full-payment due date or later. We use this classification in the

subsequent facts.[24]

Complementary evidence from bank-payment timing. The timing-based classification also predicts the payment-and-delivery pattern. Buyers classified as chained from bank-payment timing are substantially more likely to conduct transactions that are also classified as chained from invoice and delivery timing: a 71% increase relative to the mean (7.1 pp). Thus, buyers whose payments to suppliers depend most strongly on customer receipts are also much more likely to use deferred-payment arrangements in which delivery activity is first observed only when payment is due (Table C.5).[25]

Taking Stock. Together, the logistics and retail evidence show that deferred payment cannot generally be interpreted as conventional post-delivery trade credit. Across the economy, logistics activity concentrates around the end of deferred-payment windows; in retail, products purchased under deferred terms often do not enter downstream sales until after the payment deadline. These independent patterns provide unusually direct evidence that fulfillment and downstream availability are linked to settlement timing, the margin through which a liquidity disruption affects real activity in the model. The next fact asks whether the price premium on deferred transactions persists when delivery waits until the buyer is due to pay.

3.3 Fact 3. Deferred payments carry premia

In the model, for simplicity, transactions have the same nominal price, but delays reduce the effective quantity obtained through chained orders. At position n in a payment chain, a unit of expenditure yields δ^n effective units of goods, equivalently at an effective unit price $p_n^{\text{eff}} = \delta^{-n}$, relative to a spot price of one.[26] The invoice data provide a complementary perspective by measuring the monetary price per invoiced unit. We ask if this price is higher under deferred settlement, particularly for transactions whose associated delivery-service payments first appear at or after the payment due date.

[24] Additionally, in the model, longer settlement delays make chained orders more costly, creating an incentive to reduce their use. Consistent with this mechanism, sectors with longer average payment delays exhibit less chained behavior under both the bank-payment and transaction-level classifications (Appendix Figure C.3). Payment delays are measured in business days between an incoming payment and the subsequent outgoing payment among firms classified as chained; sectoral averages are weighted by sales.

[25] The comparison excludes deferred-payment transactions delivered before the due date, so the alternative is immediate payment. The result is similar when deferred transactions delivered before the due date are classified as spot rather than excluded: the difference is 5.8 percentage points, or 66% of the dependent-variable mean.

[26] Aggregating across chained orders gives $q(\mu) = \{\mathbb{E}[\delta^n \mid \text{chained}]\}^{-1}$, where the expectation uses expenditure weights.

We first ask whether the same product carries a higher unit price under deferred rather than immediate settlement. We then use the delivery-timing distinction from Fact 2 to divide deferred transactions into two groups: those for which an associated delivery-service payment is observed before the full-payment due date, as in conventional trade credit, and those for which the delivery-service payment is first observed at or after the invoice's payment deadline. This allows us to ask whether the price premium is confined to deferred transactions with logistics payments before the due date or also appears when those payments are first observed at the due date or later. Importantly, we must also consider differences due to time discounting.[27] Testing is possible in our setting because we jointly observe transaction-level prices, product identifiers, seller identities, contractual payment terms, and payments for delivery and transportation services.

Thus, for all recorded transactions in the economy, and leveraging that invoices include *unit value* prices, we estimate

$$\ln Price_{vpjt} = \delta + \gamma Deferred\,payment_{vpjt} + \lambda_p + \lambda_j + \lambda_t + \varepsilon_{vpjt}, \tag{4}$$

where $Deferred\,payment_{vpjt} = 1$ if transaction v for 13-digit product p, sold by firm j at time t, is paid through a deferred-payment invoice rather than immediately. We use this variable as the invoice-level proxy for a chained order.

Table 2 progressively narrows the comparison. Column (1) includes product, seller, and exact calendar-date fixed effects. Column (2) replaces the separate product and seller fixed effects with seller-by-product fixed effects, so γ compares immediate- and deferred-payment transactions for the same 13-digit product sold by the same seller. Column (3) further adds seller-by-date and product-by-date fixed effects, absorbing both seller-specific and product-wide price movements on the exact transaction date. Finally, Column (4) retains these demanding controls but separates deferred transactions by the *delivery timing* documented in Fact 2: whether the associated delivery-service payment is observed before the due date or first appears at or after the payment is due. Immediate-payment transactions remain the omitted category.

Table 2 shows that deferred-payment invoices carry a large price premium. In column (1), prices are 0.224 log points higher when payment is deferred, corresponding to approximately a 25% higher price. This estimate removes average price differences across products and sellers, as well as price movements common to each date. Col-

[27] For example, if a seller finances an otherwise identical invoice for D days at financing cost r, a zero-profit contract requires a deferred price proportional to $\exp(rD)$, even before accounting for any additional settlement or fulfillment cost.

Table 2: Deferred Payment, Delivery Timing, and Transaction Prices

Dependent variable: log transaction price				
	(1)	(2)	(3)	(4)
Deferred-payment invoice	0.224^{***}	0.160^{***}	0.162^{***}	
	(0.005)	(0.005)	(0.005)	
Deferred-payment: shipping at due date or after *(Chained)*				0.152^{***}
				(0.005)
Deferred-payment: shipping before payment is due *(Trade credit)*				0.170^{***}
				(0.005)
Product fixed effects	Yes	No	No	No
Seller fixed effects	Yes	No	No	No
Seller × product fixed effects	No	Yes	Yes	Yes
Seller × date fixed effects	No	No	Yes	Yes
Product × date fixed effects	No	No	Yes	Yes
Calendar-date fixed effects	Yes	Yes	No	No
Adjusted R^2	0.704	0.793	0.787	0.787
Observations	77,475,794	77,378,198	74,534,406	74,534,406
Buyer clusters	300,029	299,977	298,493	298,493

Notes: The unit of observation is a transaction. In columns (1)–(3), the independent variable is an indicator equal to one if the transaction is paid through a deferred-payment invoice rather than immediately. Column (1) includes product, seller, and exact calendar-date fixed effects. Column (2) includes seller-by-product and exact calendar-date fixed effects. Column (3) includes seller-by-product, seller-by-date, and product-by-date fixed effects, where the date interactions are defined using the exact calendar date. Common calendar-date effects in columns (3) and (4) are subsumed by the seller-by-date and product-by-date fixed effects. Column (4) retains the fixed effects in column (3) but replaces the single deferred-payment indicator with two mutually exclusive indicators based on the timing of the associated delivery or transportation-service payment. "Deferred: shipping at due date or after" identifies deferred transactions for which such a payment is first observed when full payment comes due. "Deferred: shipping before payment is due" identifies deferred transactions for which such a payment is observed before the invoice's full-payment due date. Immediate-payment transactions are the omitted category. The delivery-timing classification can be constructed only for buyers observed purchasing external delivery or transportation services. Appendix Table C.7 reports all four specifications on this common sample. Repayment windows over 360 days are excluded. Robust standard errors clustered at the buyer level are in parentheses. The data span 2024. *** denotes significance at the 1% level.

umn (2) imposes the tighter within seller–product comparison. The coefficient falls to approximately a 17% premium, but remains large and precisely estimated. Column (3) further absorbs seller-specific and product-wide price movements on the exact transaction date; the coefficient remains virtually unchanged. Thus, when *the same seller sells the same product under different payment terms*, the price remains substantially higher under deferred settlement, even after accounting for seller- and product-specific price movements over time.

Column (4) connects this price evidence directly to the delivery-timing evidence in Fact 2. Importantly for the payment-chain mechanism, transactions for which the

shipping payment is first observed at or after the payment comes due carry a premium of 0.152 log points, or approximately 16%. At a repayment horizon of roughly one month, a purely financing interpretation would imply a one-month charge of about 16%. Moreover, this estimate is nearly 90% as large as the premium for deferred transactions whose shipping payment is observed before the due date, which is the pattern more consistent with conventional trade credit.[28] The slightly higher premium when shipping payments are observed before the due date is consistent with the additional financing suppliers provide when goods are delivered before payment.

Complementary evidence from bank-payment timing. The price evidence also aligns with the independent timing-based classification. Appendix C.5, Table C.8, shows that the deferred-payment premium is larger when the buyer is classified as chained from the timing of incoming and outgoing bank payments. Thus, both the delivery-timing split and the bank-payment classification point to substantial price premia relative to spot orders in transactions that more closely resemble the payment-chain mechanism.

3.4 Fact 4. Payment chain disruptions reduce real activity

To identify the consequences of interruptions in expected payments, we investigate the effects of the intervention of Financiera Desyfin; an episode in which regulators unexpectedly froze customer accounts. Electronic invoices identify the suppliers that were awaiting payment from firms whose accounts were frozen, the amount outstanding, and the contractual date on which each payment was due. According to the model, liquidity shortages reduce firms' ability to pay upfront and, as a result, firms shift purchases from spot toward chained orders. In the language of the payment network, the freeze can potentially act like converting a funded node into a chained one, extending the sequence of payments that must clear before upstream suppliers are paid. In turn, the shift toward chained orders delays fulfillment and leaves productive resources idle, reducing TFP. Thus, a liquidity disruption can affect not only the firms directly hit but also suppliers awaiting their payments.

Credible evidence on these real effects requires both an exogenous liquidity shock and transaction-level links that identify which firms were waiting to be paid by shocked buyers. We leverage both our detailed data and the Desyfin natural experiment to study how an unexpected payment freeze propagates to exposed sellers. Namely, in August

[28] Because columns (1) and (2) use broader samples than the delivery-timing specification, Appendix Table C.7 repeats all four specifications on the column (4) sample. Estimates are nearly identical to those in the main table, showing that the column (4) findings are not driven by sample composition.

2024, Costa Rica's *Financiera Desyfin*—a financial institution catering mainly to enterprises and business clients—was abruptly and unexpectedly placed under intervention by the financial regulator, who immediately froze all customer accounts, preventing account holders from accessing their funds for months.[29] The electronic invoices identify the suppliers that were awaiting payment from those firms, the amount outstanding, and the contractual date on which each payment was due.

Outgoing transfers fall immediately at the freeze date. Before studying propagation, we first verify that the Desyfin intervention produced an immediate break in the payment behavior of firms whose accounts were frozen. Using daily Sinpe data, we restrict the sample to directly exposed firms and regress the inverse hyperbolic sine of the amount sent through interbank transfers on an indicator for dates on or after the August 13th freeze. Firm-by-day-of-week fixed effects absorb each firm's regular weekly payment pattern; the specification also includes month fixed effects. Column (1) of Table 3 shows a sharp decline at the intervention date. As a placebo, Column (2) performs an identically constructed comparison two months earlier, and yields a small and statistically insignificant coefficient. The evidence therefore shows an immediate reduction in outgoing payments at the time of the freeze, which is the payment disruption whose propagation we study next.

Propagation to upstream sellers awaiting payment. Having established that directly exposed firms reduced their outgoing payments after the freeze, we next ask whether this disruption affected sellers that were waiting to be paid by them. Our cleanest design focuses on sellers with a deferred-payment invoice outstanding at the time of the Desyfin intervention. The design combines the unexpected timing of the freeze with predetermined variation in the distance between the freeze and the invoice's pre-existing contractual due date. Conditional on being exposed to a buyer whose accounts were frozen, some sellers happened to be expecting payment within only a few days of the intervention, whereas others were not due to be paid for several weeks and potentially had time to adjust.

Let T_0 denote the Desyfin freeze date, d_i denote the contractual due date of the deferred invoice defining seller i's exposure, and $D_i \equiv d_i - T_0 > 0$ be the number of days between the Desyfin freeze and the date on which the payment was due. We estimate

[29] The regulator placed the financial institution into a resolution process under the Organic Law of the Central Bank of Costa Rica.

Table 3: Immediate Break in Outgoing Interbank Transfers at the Desyfin Freeze

Dependent variable: amount sent by firm j on business day t, asinh

	Actual freeze	Placebo date
	Aug. 13, 2024	June 13, 2024
	(1)	(2)
On or after date	-0.252^{***}	-0.020
	(0.040)	(0.036)
Month fixed effects	Yes	Yes
Firm × day-of-week FE	Yes	Yes
Observations	594,216	594,216
Firm clusters	2,358	2,358
Adjusted R^2	0.399	0.399

Notes: The table restricts the sample to firms directly exposed to Desyfin and uses daily Sinpe interbank-transfer data for 2024. The dependent variable is the inverse hyperbolic sine of the amount sent by firm j on business day t; firm-days without an outgoing transfer are coded as zero. Column (1) uses the actual intervention date, August 13, 2024. Column (2) uses the placebo date June 13, 2024. Because both specifications include month fixed effects, each coefficient is identified from the comparison before and after the corresponding date within that month. Observations from the remaining months contribute to the estimation of firm-by-day-of-week fixed effects. Column (1) therefore measures the immediate within-August break at the intervention date rather than an average effect over the entire post-intervention period. Robust standard errors clustered by firm are reported in parentheses. *** denotes significance at the 1 percent level.

$$\begin{aligned} y_{it} = {} & \beta_1 \left(D_i^{-1} \times \text{Shock}_t \right) + \beta_2 \left(\text{DesyfinLinked}_i \times \text{Shock}_t \right) \\ & + \beta_3 \left(D_i^{-1} \times \text{DesyfinLinked}_i \times \text{Shock}_t \right) + \alpha_i + \lambda_t + \varepsilon_{it}, \end{aligned} \tag{5}$$

where $\text{DesyfinLinked}_i = 1$ if, prior to the intervention, seller i had an outstanding deferred-payment invoice owed by a buyer whose funds were held at Desyfin and subsequently frozen, and $\text{Shock}_t = 1$ after the freeze. As mentioned earlier, the identifying variation is heavily disciplined: The intervention itself was abrupt and unexpected, but the specification also exploits a second, independent source of timing variation: contractual due dates were fixed before the intervention, so firms could not choose whether a promised payment would happen to be due two days or two weeks after a shock they did not foresee. The inverse term D_i^{-1} provides a parsimonious way of capturing the idea that the difference between one and two days of warning is economically more important than the difference between fourteen and fifteen days.

The fixed effects and lower-order interactions then discipline the comparison even further. Firm fixed effects α_i absorb all permanent differences across sellers, including the levels of D_i^{-1}, DesyfinLinked_i, and their interaction. Month fixed effects λ_t absorb the shock indicator itself and every change common to firms in a given month. Finally, $D_i^{-1} \times \text{Shock}_t$ absorbs any post-intervention gradient associated with due-date proximity among firms generally, and $\text{DesyfinLinked}_i \times \text{Shock}_t$ absorbs the average post-intervention change among affected sellers. Thus, β_3 is not identified from the raw

decline of affected firms. It asks whether the within-firm decline after the freeze was especially large precisely for affected sellers whose pre-existing invoices happened to be due unusually soon. We then apply the same design to sellers two links upstream: firms selling to a supplier that was itself awaiting payment from a Desyfin-affected buyer.[30]

Table 4: The Desyfin Freeze: Real Effects Along Payment Links

	TFP	Sales	Value Added
	(1)	(2)	(3)
Panel A. Sellers one link upstream (percent)			
Effect at $D = 1$	$\mathbf{-5.36}^{**}$	$\mathbf{-7.69}^{**}$	$\mathbf{-4.78}^{*}$
	(2.63)	(3.11)	(2.53)
Lower-order interactions	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Adjusted R^2	0.630	0.949	0.886
Observations	32,905	40,255	33,876
Clusters	716	794	785
Panel B. Sellers two links upstream (basis points)			
Effect at $D^{(2)} = 1$	$\mathbf{-0.127}^{***}$	$\mathbf{-0.598}^{***}$	$\mathbf{-0.800}^{***}$
	(0.018)	(0.021)	(0.029)
Lower-order interactions	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Adjusted R^2	0.585	0.928	0.847
Observations	70,357	104,266	83,641
Clusters	996	1,133	1,117

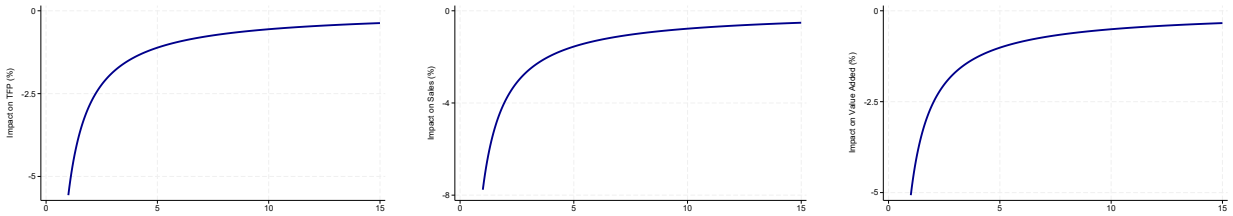
Notes: The unit of observation is a firm-month. Panel A reports $100\widehat{\beta}_3$, so the estimates are approximate percentage changes when the relevant payment was due one day after the freeze. Panel B reports $10{,}000\widehat{\beta}_3$, so the estimates are in basis points. These quantities isolate the component associated with contractual due-date proximity. At each link, D is the number of days between the freeze and the contractual due date of the payment owed to the seller at that link. Sales and value added enter in logs. TFP is estimated as described in Appendix C.6. All regressions include firm and month fixed effects and the corresponding lower-order interactions. Robust standard errors are clustered at the due-invoice buyer level. The data are monthly and span 2024. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Panel A shows that the real effects are largest when the interrupted payment was imminent. For a seller whose invoice was due to be paid by a Desyfin-affected firm one

[30]Namely, for a seller k two links upstream, let $d_k^{(2)}$ denote the contractual due date of the payment owed to k by the first-link firm, and define $D_k^{(2)} \equiv d_k^{(2)} - T_0 > 0$. Thus, at both links, the relevant timing variable is the inverse distance between the freeze and the date on which the seller at that link was due to be paid. The two-link specification replaces D_i^{-1} and $\mathrm{DesyfinLinked}_i$ in equation (5) with $(D_k^{(2)})^{-1}$ and an indicator for two-link exposure, while retaining the same lower-order interactions and fixed effects.

day after the freeze, the interrupted payment is associated with additional declines of approximately 5.4% in TFP, 7.7% in sales, and 4.8% in value added.[31]

Figure 6 plots these one-link estimates. The losses are largest for sellers whose payments were due immediately after the freeze, just as the payment-chain mechanism predicts. As the due date moves farther from the freeze, the implied losses become smaller. The intuition is that sellers expecting their money within only a few days learned, with almost no warning, that the liquidity on which they had been relying would not arrive. By contrast, sellers whose invoices were not due for several weeks learned about the freeze well before their expected payment date and therefore had more time to secure substitute liquidity, renegotiate their obligations, or otherwise adjust to the disruption. This exercise therefore provides a direct empirical analogue of the model's propagation mechanism: a shock to a buyer has the largest real effects upstream precisely when it interrupts an imminent payment, leaving the seller little time to replace the liquidity needed to pay its own suppliers and keep production moving.



(a) TFP **(b) Sales** **(c) Value Added**

Days from intervention to contractual invoice due date

Figure 6: One-Link Real Effects of the Desyfin Freeze: Due-Date Gradient

Notes: Each panel plots $100\widehat{\beta}_3/D$ from equation (5), where D is the number of days between the Desyfin freeze and the contractual invoice due date. The curves isolate the component of the post-intervention effect associated with the due date's predetermined proximity. A smaller D indicates that the seller expected the disrupted payment more imminently when the accounts were frozen.

Panel B extends the same design one transaction farther upstream. When the second-link seller's payment was due one day after the freeze, the estimated additional declines associated with due-date proximity are 0.127, 0.598, and 0.800 basis points in TFP, sales, and value added, respectively. The magnitudes are naturally smaller than at the first link, but all three estimates are negative and significant.[32] The two-link result indicates that a liquidity shock originating at Desyfin predicts real activity even two links away from the directly affected firm.

[31] We estimate monthly firm-level TFP as the residual from a Levinsohn–Petrin production function. Appendix C.6 describes the inputs and capital-stock construction.

[32] The two-link estimates are naturally small in levels because exposure compounds along the path: the product of the median pre-freeze sales-exposure shares at the first and second links is only 0.33%.

Earlier-outstanding-invoice placebo. A concern is that the main due-date gradient may reflect persistent characteristics of firms or their contracting arrangements rather than the timing of the payment obligation interrupted by the freeze. We therefore reconstruct the contractual-timing measure using the deferred invoices outstanding on June 13, 2024, two months before the intervention, and their own contractual due dates. Thus, the placebo replaces the payment obligation used to construct the timing measure with one observed in the earlier outstanding-invoice cross section, while retaining the actual Desyfin exposure and post-freeze outcome period. None of the six placebo interactions—three outcomes at each of the two links—is statistically distinguishable from zero (Appendix Table C.9). The main gradient is therefore tied to the timing of the payment obligations actually outstanding at the freeze and is not reproduced using the same firms' earlier contractual timing.

Complementary evidence from bank-payment timing. The Desyfin episode is revisited in Appendix C.8 using the timing-based classification introduced in Fact 1. Table C.10 shows that directly exposed firms became more likely to switch from spot to chained payment behavior after the freeze. It also reports declines in purchases, sales, and value added among directly exposed firms and their upstream suppliers associated with the exposed firm's realized switching into chained payments. The broader sample allows us to trace these activity patterns over time among suppliers one link upstream (Figure C.7) and across successive upstream links (Figure C.8). These results use a different data source to document how firms' payments and real activity change together, while the main analysis identifies the upstream effects of the freeze using payment obligations and contractual due dates fixed before the intervention, and yet both analysis deliver consistent results.

Taking stock. Fact 1 documents that deferred payment is widespread and is especially common among buyers whose payments to suppliers follow receipts from customers. Fact 2 supports the model's central timing assumption: goods bought with deferred payment often remain unavailable for downstream use until the buyer is due to pay. Fact 3 supports the wedge between spot and chained orders. Fact 4 provides the bridge from the transaction-level evidence to the propagation and the business-cycle implications. model in which disruptions to short-term funding shift firms from spot toward chained orders, slow settlement throughout the payment network, and depress aggregate productivity.

4 Leverage and Payment Decisions

This section endogenizes firms' choice between spot and chained orders. Section 2 showed how a the share of chained expenditures maps into TFP. Section 3 provides empirical support for the micro-economic margins: buyers with limited liquidity use deferred terms, payment deadlines predict product delivery activity, and interrupted payments reduce real activity along the chains.

We now study how short-term liquidity determines whether an agent pays upfront or waits for an earlier receipt. When short-term liquidity is scarce, agents cannot pay spot but can continue purchasing through chained orders, accepting the higher effective price implied by the payment-chain block. This section studies that decision and Section 5 closes the model in general equilibrium. For the remainder, we restrict attention to the random assignment case.

4.1 Consumption-Saving Decisions

We first study the (perfect foresight) problem of entrepreneurs who buy and sell inputs in the payments-chain network. Time is discrete $t = 0, 1, 2, \ldots$. At the start of each period, entrepreneurs receive a unit endowment of non-storable inputs available for sale and purchase inputs from others that are converted into final consumption goods. Entrepreneurs have log-utility over final consumption goods. At the start of each period, entrepreneurs place and receive spot and chained orders, as above. For this section, they take the sequence of prices $\{q_t\}_{t\geq 0}$, and interest rates, $\{R_{t+1}\}_{t\geq 0}$, as given.

At the start of t, the entrepreneur chooses *spot* purchases $S_t \geq 0$ (requiring immediate funds) and *chained* goods $X_t \geq 0$ (purchased via chained orders). Their sum, $S_t + X_t$, yields period utility, indicating that final goods are perfect substitutes. Total expenditures, $E_t = S_t + q_t X_t$, are as in the aggregate budget constraint from the previous section. Recall we treat inputs as a numeraire while $q = \mathcal{A}^{-1}$ translates chained expenditures into goods. The entrepreneur enters period t with debt B_t (negative wealth in the previous section) and carries over B_{t+1} at rate R_{t+1} to next-period debt.

We incorporate two constraints that make this problem meaningful. First, a natural borrowing limit, $B_{t+1} \leq \bar{B} \equiv 1/(1-\beta)$. Second, spot purchases require liquidity. A *spot borrowing line* (SBL) $\tilde{B}_t$ limits intraperiod liquidity:

$$S_t \leq \bar{S}_t \equiv \max\{\tilde{B}_t - B_t,\, 0\}.$$

The entrepreneur chooses spot and chained expenditures to maximize $\sum_t \beta^t \log(S_t + X_t)$. Given total expenditures E_t, since $q_t \geq 1$, the entrepreneur makes as many spot

expenditures as the SBL permits. Thus, given E_t, spot and chained expenditures are:

$$S_t = \min\left\{\bar{S}_t, E_t\right\}, \qquad X_t = \left(E_t - \min\left\{\bar{S}_t, E_t\right\}\right)/q_t. \tag{6}$$

Given this optimal expenditure split, the entrepreneur's intertemporal utility maximization can be cast into a dynamic programming problem:

Problem 1. (Entrepreneur's Recursive Problem): *Given* B_0 *and* $\left\{\tilde{B}_t, R_{t+1}, q_t\right\}_{t\geq 0}$, *the entrepreneur chooses debt holdings,* $\{B_{t+1}\}_{t\geq 0}$, *that solve:*

$$V_t(B) = \max_{B'\leq \bar{B}} \log(S+X) + \beta V_{t+1}(B')$$

where S *and* X *are given by* (6) *and total expenditures by* $E = B'R_{t+1}^{-1} + 1 - B$.

The Bellman equation expresses the entrepreneur's problem as a choice over intertemporal debt. The choice of carried-over debt, B', determines the period's expenditures, through (1), and the optimal split (6) yields S and X given the level of expenditures.[33] When short-term means of payment are scarce, agents resort to chained orders, understanding their higher cost.

The spot borrowing line (SBL), $\bar{S}_t$, represents short-term funds sourced from credit lines or trade credit, as explained earlier. These loans carry interest only if rolled over into future debt. Of course, if agents could borrow further intra-period spot funds beyond $\bar{S}$, they would as long as the intra-period loan is less than $1/q_t$. The SBL differs fundamentally from the hard debt limit, $\bar{B}$, limiting intertemporal borrowing B_{t+1}. Entrepreneurs can roll over SBL usage into future debt B_{t+1} without immediate principal repayment. This resembles a credit card limit, where beginning-of-period spending is constrained by the outstanding balance, which then carries interest forward.[34]

Characterization. We identify two debt thresholds that determine whether entrepreneurs can purchase all of their expenditures through spot orders or, conversely, whether they can make any spot expenditures at all.

Lemma 1 (Expenditure Threshold Points)**.** $S_{t+1} = 0$ *iff* $B_{t+1} \geq \tilde{B}_{t+1}$. *Define the efficiency threshold,* $B^{\star}_{t+1} \equiv R_{t+1}(\max\{\tilde{B}_t, B_t\} - 1)$. *Then,* $X_t > 0$ *iff* $B_{t+1} > B^{\star}_{t+1}$.

Next period's SBL $\tilde{B}_{t+1}$ determines if entrepreneurs can make any spot purchases in the following period. In turn, when carried-over debt exceeds the efficiency threshold $B^{\star}_{t+1}$, spending is beyond the SBL and, therefore, chained expenditures are used. These

[33] The time index in the value function reflects the evolution of $\{R_{t+1}, \tilde{B}_t, q_t\}$.

[34] Financial intermediaries often set different limits as a way to tighten short-term credit without unfeasibly demanding principal repayment, which could trigger costly defaults for both parties.

thresholds are critical for the characterization because debt levels affect the availability of spot expenditures and, thus, the price at which goods are bought at the margin, i.e., *marginal prices*, relevant for marginal decisions. The *marginal expenditure price*, $\tilde{q}_t^E(B') \equiv 1+(q_t-1)\mathbb{I}_{[B'\geq B^\star_{t+1}]}$ indicates the marginal cost of final goods when carried-over debt is B'. The *marginal borrowing price* is $\tilde{q}_{t+1}^B(B') \equiv 1+(q_{t+1}-1)\mathbb{I}_{[B'\geq \tilde{B}_{t+1}]}$ indicates the marginal cost of a good tomorrow purchased with marginal savings today.[35] Marginal prices capture the consumption-saving tradeoffs that determine the entrepreneur's optimal debt, via a generalized Euler equation:

Proposition 2 (Entrepreneur's First-Order Condition)**.** *Consider an increasing sequence of $\tilde{B}_t$. Any solution $\{B_{t+1}\}_{t\geq 0}$ to the entrepreneur's problem satisfies:*

$$\frac{E_{t+1}}{\beta E_t}\frac{Q_t}{Q_{t+1}} = \frac{R_{t+1}}{\tilde{q}_{t+1}^B(B_{t+1})/\tilde{q}_t^E(B_{t+1})} \text{ if } B_{t+1} \neq B^\star_{t+1}, \quad \beta q_t R_{t+1} \geq \frac{E_{t+1}}{E_t} \geq \beta R_{t+1} \text{ if } B_{t+1} = B^\star_{t+1},$$

where $Q_t \equiv E_t/\left(S_t + X_t\right)$ is the average prices of final goods.

Despite its unconventional format, this Euler equation retains a familiar interpretation: at interior solutions ($B_{t+1} \neq B^\star_{t+1}$), marginal rates of substitution across periods equal the return on savings corrected for price differences between spot and chained consumption. The left side represents the ratio of marginal rates of substitution (expressed as the ratio of expenditures over average prices yielding the ratio of consumption consistent with log utility), while the right side captures returns adjusted by "inflation" in terms of marginal prices. In turn, the inequalities follow when future debt is the efficient debt level, a corner solution. This Euler condition is necessary but not sufficient, as the value function is not concave: multiple sequences of future debt may satisfy the condition starting from an initial state. This non-convexity opens the door to debt-overhang leading to hysteresis in general equilibrium.

Stationary Dynamics. We now characterize the stationary dynamics of the entrepreneur's problem, i.e., when exogenous variables are constant, $\tilde{B}_t = \tilde{B}_{ss}$, $q_t = q_{ss}$, and $R_t = R_{ss} = 1/\beta$. We focus on $\tilde{B}_{ss} > 1$ and define $B^\star_{ss} \equiv B^\star(1/\beta, \tilde{B}_{ss})$.

Proposition 3 (Stationary Entrepreneur Problem)**.** *In a stationary solution: If $B_0 \leq B^\star_{ss}$, then $B_t = B_0$, $\forall t$. Moreover, there exists $B^h > \tilde{B}_{ss}$ such that: If $B_0 \in (B^\star_{ss}, B^h)$, then $B_t \to B^\star_{ss}$ in a finite number of periods. If $B_0 > B^h$, then $B_t = B_0$, $\forall t$.*[36]

[35] If carried over debt exceeds the efficiency threshold, reducing debt requires sacrificing $1/\tilde{q}_t^E$ units of consumption. If future debt exceeds the SBL, savings yield $1/\tilde{q}_{t+1}^B$ goods per unit of savings.

[36] Appendix D.1 provides a complete characterization of B^h and the number of periods needed to arrive at $B^\star_{ss}$. At B^h there is indifference between deleveraging or not.

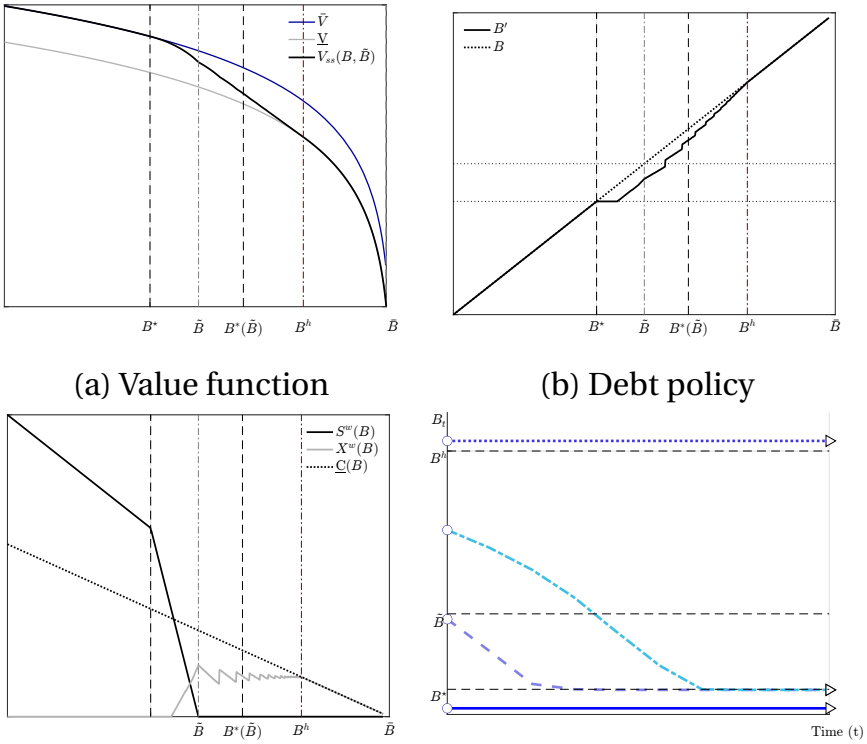
Figure 7 brings together the theoretical insights of the proposition. To illustrate the forces at play, we overlay the entrepreneur's value function with two benchmark cases: $\bar{V}$, the value function when the spot borrowing limit is unbounded, ($\tilde{B}_t = \infty$) and the entrepreneur can finance any desired level of spot expenditures, and, $\underline{V}$, where no liquidity is available ($\tilde{B}_t = 0$) and, hence, no spot expenditures are possible.

Panel (a) demonstrates how the actual value function deviates from both perfect liquidity ($\bar{V}$) and zero liquidity ($\underline{V}$) benchmarks for intermediate debt levels. For debt levels below $B^\star$, the value function V coincides with $\bar{V}$: entrepreneurs are unconstrained and consume the annuity value of their endowment net of interest payments entirely via spot purchases, showcasing that for sufficiently low debt levels, the SBL is irrelevant. When debt exceeds the high threshold B^h, the value function lies on $\underline{V}$, indicating that entrepreneurs consume exclusively via chained purchases forever. Crucially, although a highly indebted entrepreneur could save to secure better future prices, he does not, as the gains from deleveraging arrive too far in the future.

Consistent with Panel (a), Panel (b) shows non-monotonicity in the optimal debt policy function. Entrepreneurs with debt below $B^\star_{ss}$ have constant debt levels; entrepreneurs with debt above B^h remain permanently at their initial debt levels. Only those with moderate debt in $(B^\star_{ss}, B^h)$ gradually deleverage toward the efficient threshold $B^\star$ by choosing future debt (solid line) below current debt (dashed line). Panel (c) shows how consumption policies differ across debt levels. When debt is below (above) $B^\star_{ss}$ (B^h), the entrepreneur consumes the annuity of his wealth as spot (chained). For debt levels slightly below B^h, we observe a reduction in overall consumption to deleverage. The chain-saw pattern of the optimal policy reflects a trade-off between smoothing consumption and shortening the number of periods to arrive at the efficient debt level. Panel (d) traces the debt dynamics that emerge from the non-monotonic policy functions: Starting from different initial debt levels, entrepreneurs either remain at their initial debt levels or temporarily delever toward the efficient debt threshold.

The main takeaway is that, unlike canonical financial frictions models, entrepreneurs lack incentives to delever, even when their discount factor equals the real interest rate. Indeed, if sufficiently indebted, they prefer to stay in debt and pay the high price of chained orders. This distinction occurs as the benefits of deleveraging do not happen at the margin. On the contrary, the partitioning of the state space into two regions with different domains of attraction is a shared feature with models where, due to financial constraints, there are rewards to bringing debt to a certain level. Whereas the rewards here stem from accessing better prices, in other models rewards stem from avoiding

roll-over risk (Cole and Kehoe, 2000; Aguiar et al., 2015) or technologies requiring scale (Buera and Shin, 2013).[37] The lack of incentives to delever has significant consequences for the general equilibrium that are aggravated by the pecuniary externalities.



(a) Value function (b) Debt policy

(c) Consumption policy (d) Transitions

Figure 7: Bellman Equation and Policy Functions: Entrepreneur's Problem

Notes: Figures computed using value function iteration with $\beta = 0.8$, $R = \beta^{-1}$, $q = 1.75$, and $\tilde{B} = 0.4 \cdot \bar{B}$. $\bar{V}$ and $\underline{V}$ are value functions with an unlimited and a zero SBL. $\underline{V}$ is the entrepreneur's consumption associated with the value function $\underline{V}$.

4.2 Investment Decisions

We can also study how payment chains can affect investment and production by incorporating these features into the entrepreneur's problem. For that, we study a problem where both consumption and investment must be purchased in the payment network and part of yesterday's purchases are inputs for today's production.

Problem 2. (Entrepreneur's Problem with Production): *Given* B_0 *and* $\left\{\tilde{B}_t, R_{t+1}, q_t\right\}_{t\geq 0}$, *the entrepreneur chooses debt-holdings and investment* $\{B_{t+1}, K_{t+1}\}_{t\geq 0}$, *the solution to:*

$$V_t(B,K) = \max_{B'\leq \bar{B}, K'\geq 0} \log(S + X - K') + \beta V_{t+1}(B', K')$$

[37] A continuous-time limit akin to those studied in roll-over risk models shows that the saw-like pattern of numerical solution disappears and that the value functions are isomorphic—the saw like pattern appears in discrete time due to the integer periods of arrival to unconstrained debt position; see the online note, "Consumption–Saving Problem with Wealth Rewards."

where S and X are given by (6) *and total expenditures by* $E = B'R_{t+1}^{-1} + f(K) - B$.

Instead of an endowment, the entrepreneur has $f(K)$ goods produced with input (e.g., capital or inventory) K acquired in the previous period. In turn, currently purchased goods, $S + X$, are distributed between consumption, C and inputs for next-period production, K', assuming full depreciation to keep the analysis as simple as possible.

Recall that in the endowment problem above, the entrepreneur's stationary problem reaches two types of steady-states: either all-spot-expenditures or all-chained-expenditures steady states. The same two types of steady states are reached in the problem with investment. Furthermore, away from steady state, the same incentive to reduce debt and to confront better prices for moderate debt levels is present in this problem, leading to types of steady states and an intermediate region, a domain of attraction toward the efficient steady state.

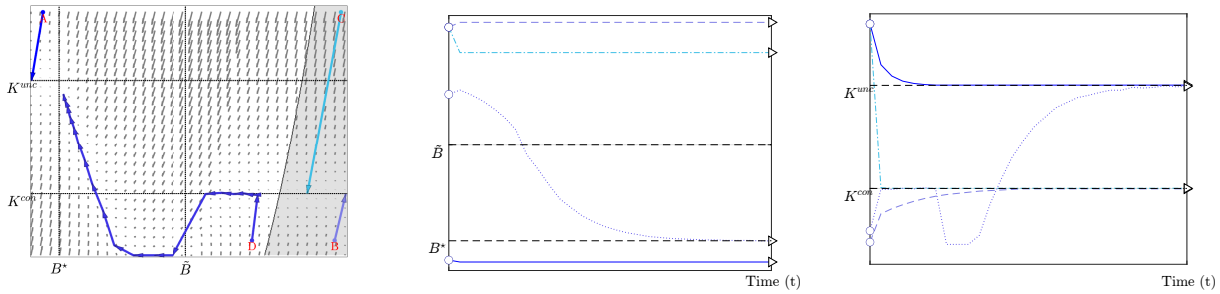
The different insight emerging from the version with production is that the higher cost of chained goods also distorts the optimal investment scale. If the entrepreneur only makes spot purchases, the return on investment equals the rate of interest: $f'(K_{ss}) = R = \frac{1}{\beta}$. However, when all expenditures are chained, the first-order condition for inputs becomes: $f'(K_{ss}) = qR = \frac{q}{\beta}$. Hence, the marginal product of capital exceeds the rate of return on saving and investment, as in other distorted economies. In intermediate cases, investment is distorted by average prices.

Figure 8 illustrates the domains of attraction toward the two classes of steady states and the rich dynamics that emerge with production. Panel (a) shows the phase portrait[38] with arrows indicating the direction of change for both state variables (B, K) of the model's solution together with some trajectories. The shaded region represents the domain of attraction toward the disrupted steady state with only chained expenditures and inefficient investment. Four representative paths demonstrate the range of possible dynamics. Path A begins with low debt and excess capital: convergence to the efficient steady state quickly as excess income is used to reduce debt, reaching equilibrium with undistorted investment while maintaining the pattern of only spot expenditures. Paths B and C start with high initial debt but different capital levels—both converging to the steady state with low investment and only chained expenditures. The most interesting case, Path D, starts with intermediate debt and low capital. Initially, the entrepreneur accumulates both debt to increase capital and reach an efficient scale. It begins to payoff debt and once debt falls sufficiently to enable some spot purchases,

[38] A phase portrait draws the discrete vector field, where arrows indicate the direction and magnitude of change from each point, one (B, K) to (B', K').

investment is reduced to accelerate the deleveraging until the efficient level of capital is reached and the entrepreneur makes only spot expenditures.

Panels (b) and (c) trace the corresponding debt and capital paths over time, confirming that the same two-steady-state pattern holds with production: entrepreneurs either converge quickly to efficiency or remain trapped in a state of persistent underinvestment. The same pattern of debt overhang, when debt starts high but deleveraging when debt is low, is observed in panel (b). Panel (c) shows how capital evolves depending on the marginal prices.



(a) Phase Portrait **(b) Path of Debt** **(c) Path of Capital**

Figure 8: Investment and Leverage Decisions

Note: Figures are calculated using value function iteration with $\beta = 0.8$, production function $f(k) = k^{\alpha}$ with $\alpha = 0.5$, $q = 1.25$, $R = \beta^{-1}$, and $\tilde{B} = 0.2 \cdot \bar{B}$. Panel (a) corresponds to the phase-space diagram. Panels (b) and (c) correspond to the transitions for debt and capital, respectively.

5 General Equilibrium

The previous section analyzes how entrepreneurs choose between spot and chained purchases given exogenous terms. We now embed these decisions into a general equilibrium where prices and interest rates are endogenous, revealing how individual payment choices create aggregate externalities. We provide insights about crises.

Environment. We provide a simple setting with endogenous savings and expenditures decisions. Timing remains identical to the partial equilibrium problems above: payments are executed within the period and debt is carried over. As in Kiyotaki and Moore (1997b), the economy consists of two groups of agents. *Entrepreneurs* are the same agents studied in the previous section, but we specialize their aggregate input endowment to $y = 1$. The new agents are *creditors* who hold positive financial wealth but receive no endowment. Both share the same log preference. To consume, creditors always carry positive wealth, D_t, earning the gross return R_{t+1} between periods. Credi-

tors serve as the economy's savers, providing the funds that entrepreneurs borrow.[39]

Problem 3 (Creditor's (Recursive) Problem)**.** *Given D_0 and $\{R_{t+1}\}_{t\geq 0}$, the creditor chooses deposit holdings $\{D_{t+1}\}_{t\geq 0}$ to solve:*

$$V_t(D) = \max_{D'\geq 0} \log(C) + \beta V_{t+1}(D')$$

subject to the intertemporal budget constraint: $C + R_{t+1}^{-1} D' = D$.

The solution is standard to logarithmic preferences: $C_t = (1-\beta)D_t$ and $D_{t+1} = \beta R_{t+1} D_t$ for all $t \geq 0$. Their total expenditure each period is $C_t = (1-\beta)D_t$. Implicitly, we use that creditors always have positive wealth, as they only make spot expenditures.[40]

Equilibrium. Since entrepreneurs supply the input endowment, $y = 1$, which is the numeraire, *clearing in the input market* requires total expenditures to equal the input endowment, $E_t^s + E_t^x = 1$. *Clearing in the asset market* requires the creditors' deposits exactly equal entrepreneurs' debt: $D_t = B_t$. Spot expenditures are the sum of the creditors' consumption (who only make spot purchases) and entrepreneurs' spot purchases: $E_t^s = C_t + S_t$. Chained expenditures come only from entrepreneurs: $E_t^x = q_t X_t$. From Proposition 1, chained expenditures are $\mu_t = q_t X_t$, which determines *a price consistency condition* $q_t = \mathcal{A}^{-1}(\mu_t)$. Because spot expenditures translate one for one to output we have that *final goods consumption equals total production* $C_t + S_t + X_t = \mathcal{Y}(q_t X_t)$.

Definition 2 (Symmetric Competitive Equilibrium)**.** *Given a sequence of borrowing limits $\{\tilde{B}_t\}_{t\geq 0}$, sequence $\{B_t, C_t, S_t, X_t, R_t, q_t\}_{t\geq 0}$ is a* ***symmetric competitive equilibrium*** *it satisfies $\forall t$:* ***I. Optimization*** *(given $\{R_t, q_t\}_{t\geq 0}$, the sequences $\{D_t, B_t, S_t, X_t\}_{t\geq 0}$ solve the entrepreneur's problem and $\{D_t, C_t\}_{t\geq 0}$ the creditor's problem),* ***II. Asset and input markets clear****,* ***III. Final goods consumption equals production****, and* ***IV. Price consistency****.*

The equilibrium definition is standard, except that q_t and $\mathcal{Y}_t$ depend on the amount of chained expenditures. Despite the richness of network features, the economy operates as an economy with production externalities on $\mathcal{Y}$ and pecuniary externalities on q, encoding the payments technology and delays in the payments-chain network.

Equilibrium Dynamics: Recursive Representation. A virtue of this setting is that it renders a tractable way to characterize the equilibrium recursively. Given that $D_t = B_t$,

[39] Creditors play a key role: they close the asset market and their spot purchases guarantee a mass of spot expenditures that prevents the full economy from collapsing if entrepreneurs' debt exceeds the SBL.

[40] That is, their budget constraint also takes the form of (2) with $X = 0$.

B_t is the single endogenous state. The state variable is the vector $\{B_t, \tilde{B}_t, \tilde{B}_{t+1}\}$. A recursive equilibrium is a mapping from the state to the vector of prices R_{t+1}, q_t, and allocations, $\{C_t, S_t, X_t\}$, and future endogenous state B_{t+1}. Because a recursive representation is possible, we drop time subscripts, adopting standard recursive notation. Since only B is endogenous, we seek a mapping $\mathcal{B}$ that determines tomorrow's debt, B', as a function of today's state: $B' = \mathcal{B}(B, \tilde{B}, \tilde{B}')$.

We begin by establishing that the time t allocation is only a function of B and $\tilde{B}$.

Proposition 4 (Equilibrium Allocation given $\{B, \tilde{B}\}$)**.** *Given debt B and spot borrowing limit $\tilde{B}$, the equilibrium allocation is characterized by:*

1. ***Creditor consumption**: $C(B) = (1-\beta)B$ (creditor log expenditures)*
2. ***Entrepreneur expenditures**: $E(B) = 1 - (1-\beta)B$ (input-market clearing)*
3. ***Spot purchases**: $S(B, \tilde{B}) = \min\{\max\{0, \tilde{B} - B\}, E(B)\}$ (spot maximization)*
4. ***Chained expenditures**: $\mu(B, \tilde{B}) = E(B) - S(B, \tilde{B})$ (definition of total expenditures)*
5. ***Network price**: $q(B, \tilde{B}) = \mathcal{A}^{-1}(\mu(B, \tilde{B}))$ (definition)*
6. ***Chained goods**: $X(B, \tilde{B}) = \mu(B, \tilde{B})/q(B, \tilde{B})$ (definition)*
7. ***Average price**: $Q(B, \tilde{B}) = E(B)/(S(B, \tilde{B}) + X(B, \tilde{B}))$ (definition).*

This proposition shows that the current debt level and SBL yield the entire allocation of goods. What remains is to characterize the evolution of debt, i.e., finding $\mathcal{B}$. This debt mapping emerges from combining the entrepreneur's and creditor's Euler equations. From the creditor's Euler equation, the gross real interest rate can be expressed as a function of B and B'; namely, $R(B, B') = B'/(\beta B)$. Thus, we can replace this gross rate in the entrepreneur's interior Euler equation, Proposition 2, to obtain

$$\underbrace{\frac{B}{E(B)} \cdot Q(B, \tilde{B})}_{\equiv \mathcal{E}(B; \tilde{B})} = \underbrace{\frac{B'}{E(B')} \cdot Q(B', \tilde{B}') \cdot \frac{q^E(B'; B, \tilde{B}, \tilde{B}')}{q^B(B'; \tilde{B}')}}_{\equiv \mathcal{E}'(B'; B, \tilde{B}, \tilde{B}')} \tag{7}$$

where $q^E(B'; B, \tilde{B}, \tilde{B}') \equiv 1 + (q(B, \tilde{B}) - 1)\mathbb{I}_{[B' \geq B^\star(\tilde{B})]}$ and $q^B(B'; \tilde{B}') \equiv 1 + (q(B', \tilde{B}') - 1)\mathbb{I}_{[B' > \tilde{B}']}$ are the marginal expenditure and borrowing prices as functions of the state.

When interior solutions to the entrepreneur's problem are valid, the debt mapping is determined by solving $\mathcal{E}(B; \tilde{B}) = \mathcal{E}'(B'; B, \tilde{B}, \tilde{B}')$ for B'. Due to the discontinuous nature of marginal prices, this equation may have multiple solutions, as in partial equilibrium. However, the partial equilibrium problem tells us how to select the appropriate solution, so long as the entrepreneur wants to delever at given prices:

$$\mathcal{B}(B, \tilde{B}, \tilde{B}') = \max\left\{B^\star(\tilde{B}), \arg\min_{B'}\{\mathcal{E}(B; \tilde{B}) = \mathcal{E}'(B'; B, \tilde{B}, \tilde{B}')\}\right\}.$$

This mapping ensures entrepreneurs never choose debt levels below the threshold $B^\star(\tilde{B})$ where they would be unconstrained, and among multiple solutions, they select the one that keeps continuity in their value function. It is a non-linear implicit difference equation that provides a recipe to solve the model: first, solve for the equilibrium allocation and price functions as a function of $\{B, \tilde{B}\}$. Then, implicitly solve for $\mathcal{B}$. We exploit this feature to obtain a graphical device to analyze equilibria and several properties.

5.1 Implication I: Transitional Dynamics and Hysteresis

Steady States. In a steady state, the SBL, $\tilde{B}_{ss}$, and debt must be constant: $B_{ss} = B$. From the creditor's solution, a necessary and sufficient condition for any steady state is that the real interest rate is constant at $R_{ss} = \beta^{-1}$. Thus, a specific debt level constitutes a steady state if, given this interest rate and the equilibrium price q, entrepreneurs optimally choose to keep that same debt, a condition we studied in partial equilibrium.

Since at a steady state with $R_{ss} = 1/\beta$ entrepreneurs either only make spot or only make chained expenditures, we conclude there are two types of steady state in general equilibrium—there is a continuum of steady states depending on B_0. In an *undisrupted steady state*, debt remains below $B^\star$, entrepreneurs exclusively make spot purchases, $q = 1$, payments are immediate, and the economy operates at full potential with $\mathcal{Y}_{ss} = 1$. In a *disrupted steady state*, entrepreneurs rely entirely on chained purchases, resulting in $q > 1$, delayed payments, and output below potential: $\mathcal{Y}_{ss} < 1$.

Corollary 1 (Steady-State Characterization)**.** *Fix $\tilde{B}_t = \tilde{B}_{ss} > 0$. Then, for all t:*

i. The economy is in an undisrupted steady state when $B_t \leq B^\star(\beta^{-1}, \tilde{B}_{ss})$.

ii. The economy is in a disrupted steady state when $B_t \geq B^h(\beta^{-1}, \tilde{B}_{ss})$.

The disrupted steady state is an inefficient hysteresis region where excessive debt permanently impairs the economy's payments efficiency. The economy is constrained inefficient because agents do not internalize how their expenditures affect q.[41]

Domain of Attraction Toward Efficient Steady States. For debt levels between the two steady-state regions, the economy may experience transitional dynamics toward the undisrupted steady-state region. The domain of attraction toward undisrupted steady states depends on whether the aggregate Euler equation $\mathcal{E}(B; \tilde{B}) = \mathcal{E}'(B'; B, \tilde{B}, \tilde{B}')$ has a solution with $B' < B$—that is, entrepreneurs choose to delever. This domain of attraction toward undisrupted steady states has an upper bound,[42] B^* given by:

[41] This mechanism provides a theoretical foundation for understanding persistent periods of low productivity among highly indebted firms observed during extended economic stagnation.

[42] For debt levels exceeding this upper bound, the economy does not transition toward the efficient steady state in a symmetric equilibrium.

$$B^*(\tilde{B}) \equiv \tilde{B}/(S(\tilde{B},\tilde{B}) + X(\tilde{B},\tilde{B}) + C(\tilde{B})) \geq \tilde{B}.$$

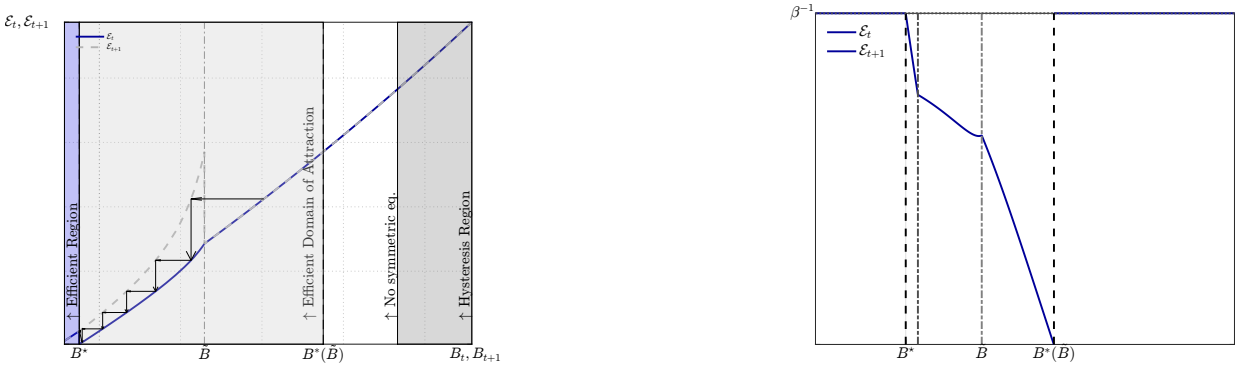
Corollary 2 (Convergence to Efficiency)**.** *Let $\tilde{B}_t = \tilde{B}_{ss}$. For any $B_0 < B^*(\tilde{B}_{ss})$:*

1. *If $B_t \in (B^\star(\beta^{-1}, \tilde{B}_{ss}), B^*(\tilde{B}_{ss}))$, then $B_{t+1} < B_t$ (deleveraging occurs)*
2. *If $B_t \leq B^\star(\beta^{-1}, \tilde{B}_{ss})$, then $B_{t+1} = B_t$ (steady state reached).*

Payments-chain delays are only *temporary* when the SBL is at steady state, provided that debt starts in the domain of attraction toward efficient steady states.

Figure 9 illustrates the dynamics encoded in $\mathcal{B}$. Panel (a) plots the functions $\mathcal{E}$ (solid blue) and $\mathcal{E}'$ (dashed gray) against debt levels, holding $\tilde{B}$ fixed. The arrows trace the equilibrium path of debt: when the functions intersect below the 45-degree line, entrepreneurs deleverage ($B' < B$). The convergence region extends from $B^\star$ to B^*, while the gray shaded area represents the hysteresis region. Panel (b) shows how equilibrium interest rates vary across debt levels, remaining at β^{-1} in both steady-state regions.

Finally, we note that the domains of attraction toward efficient steady states and the hysteresis region are disconnected. When debt falls in the intermediate range $\{B^*, B^h\}$, symmetric competitive equilibria do not exist. In this middle region, entrepreneurs individually would want to deleverage, but in general equilibrium, this would require interest rates far below β^{-1}, for which no solutions exist to the equilibrium condition. We leave the analysis of asymmetric equilibria for future work.



(a) Transition plot **(b) Equilibrium rates**

Figure 9: Equilibrium Phase Diagram and Hysteresis Region

Note: Figures are calculated using value function iteration with $\beta = \delta = 0.95$, and $\tilde{B} = 0.2 \cdot \bar{B}$. Panel (a) shows the phase diagram of B_t constructed using $\mathcal{E}(B'; B)$ and $\mathcal{E}'(B'; B, \tilde{B}, \tilde{B}')$. Arrows indicate the projection from one point in the Euler equation to another. The projection to the x-axis traces that path of debt. Panel (b) plots $R(B, \tilde{B}, \tilde{B}')$ over $[B^\star - 1,\ B^* + 2]$. The shaded area corresponds to the hysteresis region.

5.2 Implication II: Credit Crunches and a Resilience Paradox

Credit Crunches. We now analyze the economy's response to temporary disruptions in short-term funding, i.e., a mean-reverting tightening of borrowing limits $\{\tilde{B}_t\}$.[43] As long as initial debt satisfies $B_0 \leq B^*(\tilde{B}_{ss})$ and $\tilde{B}_{ss} > 1$, the economy eventually reverts to an undisrupted steady state. While the exercise is not meant to be a comprehensive quantitative analysis, it is calibrated to represent a full shutdown of credit, consistent with the Desyfin experiment in Section 3 but extended to all the entrepreneurial sector. Recall the model features a minimalist set of parameters, $\tilde{B}_{ss}, \delta, \beta$. We fix δ at 0.995, indicating a loss of 0.5% of the real value of goods for each additional position in the chain. This is consistent with evidence in Section 3 showing that repayment windows predict delivery timing and that chained orders are costly.[44]

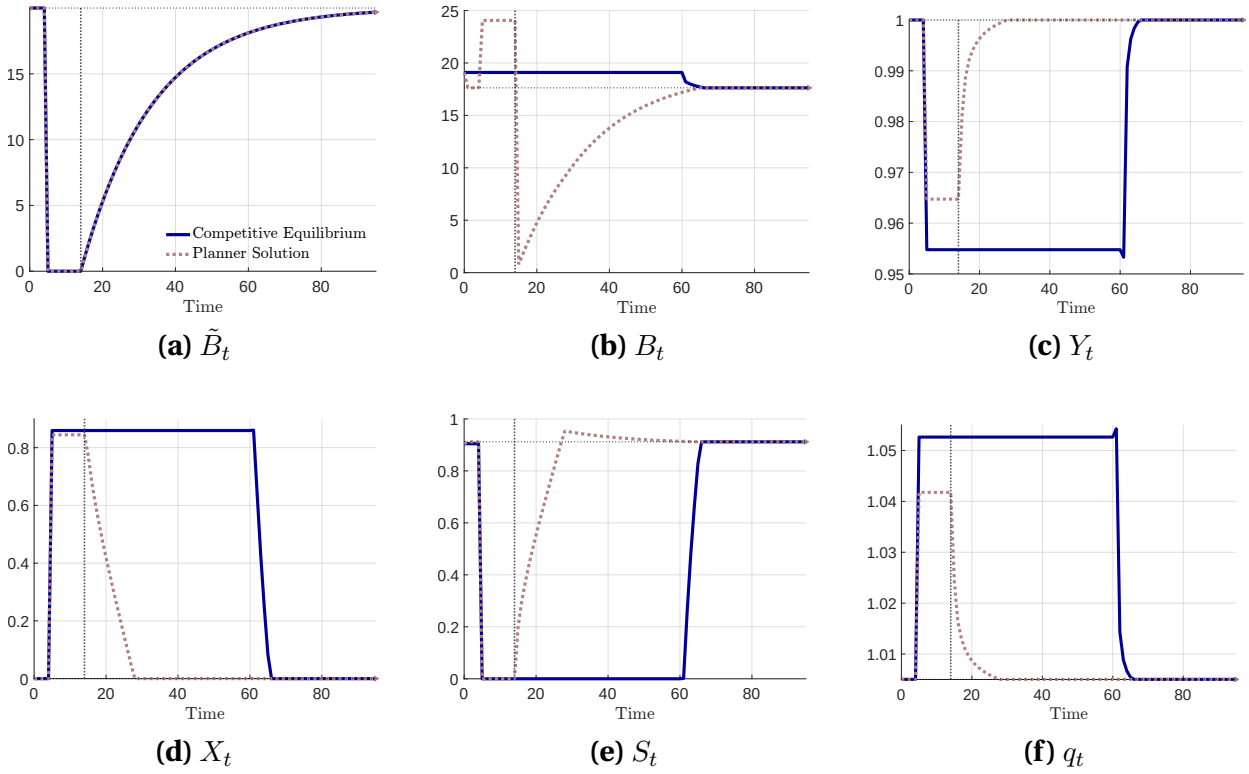
Figure 10 illustrates a transition following a credit crunch starting from an undisrupted steady state. The solid line shows the competitive equilibrium response, the dashed line represents the constrained planner's solution, discussed below. The transition exhibits three distinct phases: The *extreme phase* occurs when $\tilde{B}_t$ falls to zero, fully eliminating access to short-term credit so that no spot expenditures are feasible by entrepreneurs. On impact, the value of expenditures by creditors and entrepreneurs remains the same, but entrepreneurs switch entirely to chained expenditures. This leads to an increase in average prices and a decline in real GDP of 4.5%. The extreme phase is followed by a *hysteresis phase* where some spot expenditures would be possible if debt were reduced. In this phase, debt remains constant because entrepreneurs keep chained expenditures constant, and TFP remains depressed at its lowest value even though the SBL is recovering. Finally, there is a *deleveraging phase* where debt declines, TFP recovers quickly, and full potential is restored.

The most interesting aspect of this transition is the contrast between the hysteresis and deleveraging phases. During the deleveraging phase, entrepreneurs anticipate that prices will improve the following period because, for the current level of B, the recovery of the SBL will allow some spot expenditures on the margin. Thus, it is worth sacrificing chained expenditures for spot expenditures. This behavior reduces the level of debt, lowers the interest rate, and generates a positive externality, which is why output recovers quickly. In contrast, the hysteresis region is characterized by depressed TFP and no willingness to delever. On the margin, entrepreneurs understand that reducing chained expenditures today would only enable more chained expenditures the following period, with no access to spot expenditures. Thus, they keep their debt constant.

[43] Debt mapping $\mathcal{B}$ characterizes the full transition path, by running forward B's difference equation.

[44] We set the steady-state SBL, $\tilde{B}_{ss}$, to 10% of the natural borrowing limit. We set β to 0.995, consistent with a 6% real annual rate if the time period is a month.

The planner's solution follows a completely different path, as we discuss below.



(a) $\tilde{B}_t$ **(b)** B_t **(c)** Y_t

(d) X_t **(e)** S_t **(f)** q_t

Figure 10: Efficient and Competitive Transition after a Credit Crunch

Note: This figure reports a numerical example of a credit-crunch episode. Figures are calculated using value function iteration with $\beta = \delta = 0.995$, and $\tilde{B}_{ss} = 0.1 \cdot \bar{B}$. The figure is computed by setting 5 periods at an initial steady state, followed by a credit crunch where $\tilde{B}_t$ is set to zero for 10 periods. From $t = 15$, $\tilde{B}_t = 0$ on, B_t returns to steady state according to an AR(1) process with coefficient 0.95. The planner solution uses the Pareto weights consistent with the terminal steady state, $\tilde{B}_t = \tilde{B}_{ss}$.

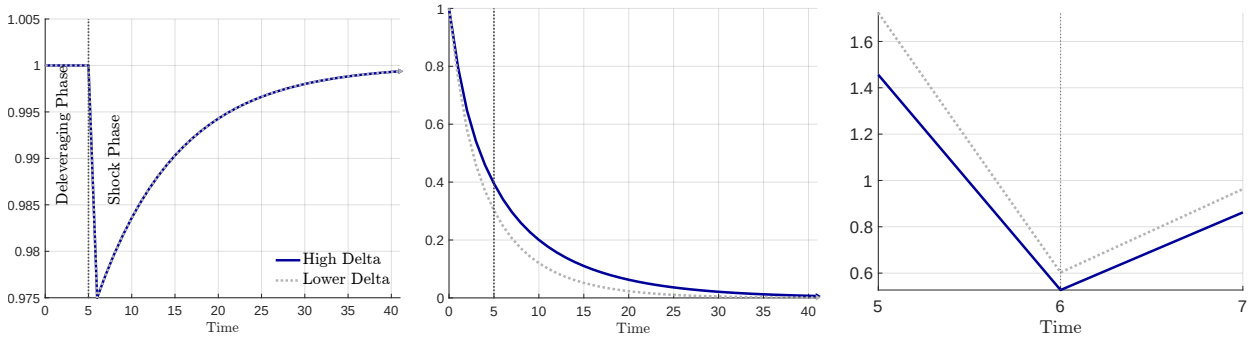
A Resilience Paradox. A natural question is whether economies with faster payments are always more resilient to credit crunches. The answer is no, given a resilience paradox: economies with lower δ (greater delays) may exhibit superior resilience to credit shocks, despite always experiencing larger efficiency losses given μ—recall Figure 3c.

The mechanism behind this resilience paradox operates through different deleveraging incentives before a shock. A lower δ means longer payment delays in the payments-chain network, making chained purchases relatively more costly. Thus, greater costs create stronger incentives for entrepreneurs to delever. A lower δ then induces a fundamental trade-off: greater payment inefficiency in the short run, given a level of chained expenditures, but reduced vulnerability to credit crunches through incentives to delever.

Figure 11 shows the paradox through a calibrated comparison. Consider two economies that begin with identical debt levels but differ in their payment efficiency parameter δ. After allowing sufficient time for each economy to reach its preferred debt level, we

subject both to an identical credit shock—the reduction in $\tilde{B}_t$ shown in Panel (a).[45]

Panel (b) reveals the key mechanism: when the shock hits, the low-δ economy (solid line) has achieved significantly lower debt levels than the high-δ economy (dashed line) due to stronger deleveraging incentives from costlier chained purchases. Panel (c) shows the paradox: despite being less efficient at processing payments, the low-δ economy experiences a smaller output decline during the credit crunch. The credit crunch occurs at a time when debt is lower in the more inefficient economy, which reduced debt faster to avoid the reliance on more expensive chained purchases.



(a) $\tilde{B}_t$ **(b)** B_t **(c)** $\Delta Y_t\%$

Figure 11: Resilience Comparison: Different values of δ.

Notes: The figure is a numerical example of a credit crunch episode comparing two economies with different production discount factors ($\delta = 0.90$ vs. $\delta = 0.855$). Figures are computed via value-function iteration: $\beta = 0.90$, and $\tilde{B}_{ss} = 0.1 \cdot \bar{B}$. The economy has an initial debt of $B_0 = \tilde{B}_{ss}$ (outside steady state). The crunch occurs in period t=6, when $\tilde{B}_t$ drops to $0.975 \cdot \tilde{B}_{ss}$. Then $\tilde{B}_t$ returns to steady state according to an AR(1) process with coefficient 0.9. The transition ends at t=41 when $\tilde{B}_t = \tilde{B}_{ss}$.

The resilience paradox requires some qualifications. (1) The paradox only manifests when both economies are in a transition toward an efficient steady-state, and we let enough time so that the low-δ exhibits less debt. The low-δ deleverages faster, but needs time to become more resilient. Indeed, if the shock happens and both economies have the same level of debt, the low-δ economy will always do worse. (2) For extremely severe shocks, e.g., $\tilde{B}_t = 0$, the low-δ economy will necessarily perform worse.

Regardless, the resilience paradox reveals an economic insight: payment system inefficiencies, while costly during normal times, can induce deleveraging that builds shock resilience. The empirical analysis is suggestive of this deleveraging-to-avoid-costs pattern: sectors with longer average delays exhibit less chained behavior (Figure C.3). The takeaway is that economies with more efficient payment systems may inefficiently maintain higher debt levels, leaving them vulnerable when credit conditions deteriorate. Next, we make precise statements about inefficiency.

[45] The shock is anticipated, but the same is true if it is not.

5.3 Implication III: Constrained Inefficiency

In Section 2, we showed that an inefficient organization of payments can lead to productive inefficiencies. In practice, governments may not be able to directly reorganize payments. However, they can influence expenditures via policy interventions.[46] We now study a planner problem that respects the payments technology and financial constraints, but chooses the evolution of debt. The analysis isolates the expenditure externalities that differ from the organizational inefficiencies of the payments technology.

The Primal Planner Problem. We focus on constrained inefficiency during credit crunch episodes as the one above. Recall a credit crunch starts from an undisrupted steady state and ends in another one. Since the economy is Pareto and productively efficient in an undisturbed steady state, Pareto inefficiencies can arise, if at all, only during a transition path—*although a distorted steady-state can also be constrained inefficient.* To study constrained inefficiencies within a transition, we consider a social planner with assigned Pareto weights that guarantee the same allocation in the terminal undisrupted steady state. This guarantees that the planner does not have redistributive objectives in the long run that would cause its chosen path of debt to deviate from that of the market outcome. This benchmark requires setting Pareto weights on entrepreneurs, θ, so that

$$\frac{S_{ss}}{C_{ss}} = \frac{1-(1-\beta)\,B_{ss}}{(1-\beta)\,B_{ss}} = \frac{\theta}{1-\theta}, \tag{8}$$

where B_{ss} is the terminal market equilibrium steady-state debt after the credit crunch.

We analyze a primal planner problem who directly chooses debt sequences $\{B_t\}_{t\geq 0}$ to maximize welfare, respecting the transactions technology and payment-chain frictions. Any allocation in this problem can be implemented with an appropriate sequence of debt and entrepreneur income taxes—see Appendix F.4.

Problem 4. (Primal Planner Problem): *Taking* $\left\{\tilde{B}_t\right\}$ *as given:*

$$\max_{\{B_t\}_{t\geq 0}} \sum_{t\geq 0} \beta^t \mathcal{P}\left(B_t, \tilde{B}_t\right),$$

$$\textit{where } \mathcal{P}\left(B,\tilde{B}\right) \equiv (1-\theta)\log\left((1-\beta)B\right) + \theta \log\left(S\left(B,\tilde{B}\right) + X\left(B,\tilde{B}\right)\right), \tag{9}$$

and taking as given $S\left(B,\tilde{B}\right)$ *and* $X\left(B,\tilde{B}\right)$ *as defined by Proposition 4.*

For a given debt, by using the conditions in Proposition 4, the constrained planner respects the agents' expenditure decisions and input market-clearing conditions.

[46] Recall that relative to a random assignment, if it could, a planner would reorganize payments so that each chain is of equal length. Traffic regulation offers an analogy: governments cannot assign drivers to fast and slow lanes, but they can tax vehicle purchases to influence traffic.

Hence, the constrained planner respects the payment-chain technology (and production resource constraint) indirectly through the functions $S(B, \tilde{B})$ and $X(B, \tilde{B})$. As the objective is time-separable[47] the optimal debt at t depends only on the current SBL $\tilde{B}_t$:

$$B_t = B^p(\tilde{B}_t) = \arg \max_{B \in [0, \bar{B}]} \mathcal{P}\left(B, \tilde{B}_t\right).$$

The solution is characterized in Appendix F. We summarize a key insight here. The solution reveals that during credit crunches, debt or its counterpart, interest rates, may be either inefficiently high or low. The reason is that the planner has two ways to induce more spot expenditures to increase TFP: it can redistribute wealth toward entrepreneurs (reducing their debt to free up spot borrowing) or toward creditors (increasing debt to raise their wealth) to stimulate their spot expenditures.

This ambivalent nature leads to four qualitatively different regimes depending on the tightness of the SBL, $\tilde{B}$. When $\tilde{B}$ is sufficiently ample, the economy operates efficiently with $\mathcal{Y} = 1$. This is possible only if $\tilde{B}$ is sufficiently large so that at $B = B_{ss}$, spot expenditures $S(B, \tilde{B}) = 1 - (1-\beta)B$ are feasible. As $\tilde{B}$ falls below that efficiency threshold, the planner faces a trade-off between productive efficiency and redistribution of resources toward entrepreneurs to increase their relative consumption in line with the Pareto weights. In moderate-constraint regions, the planner fully sacrifices redistribution to maintain the productive efficiency goal.[48] It does so by reducing debt, redistributing wealth one for one with $\tilde{B}$, guaranteeing that the entrepreneur only spends spots, but distorting its unconstrained consumption distribution, dictated by (8).

As the SBL falls further, the planner accepts some productive inefficiency, permitting chained expenditures but still lowering debt via the reduction in $\tilde{B}$, again redistributing from creditors toward entrepreneurs, balancing productive efficiency gains against worse redistribution. Strikingly, when constraints become extreme, the planner reverses strategy entirely: rather than easing the debt burden of constrained entrepreneurs, the planner redistributes wealth toward the *unconstrained agent*, the creditor. It essentially gives up on achieving productive efficiency by reducing debt and reverses course. This seemingly paradoxical result occurs because marginal increases in creditor wealth generate more spot expenditures which increases efficiency. The ambivalent nature distinguishes payment-chain crises from macro-finance models, where redistributive and productive efficiency typically reinforce each other. Here, the plan-

[47] The objective is time-separable because current debt determines creditor expenditures through log utility, with entrepreneur expenditures following residually.

[48] In typical planner problems, the objective is smooth. Hence, the planner typically balances objectives. Here, the planner forgoes redistribution of objectives, a corner solution that results from the discontinuity of marginal TFP losses from chained expenditures.

ner may switch between complementary and conflicting policies as constraint severity varies because productive efficiency can be increased by redistributing wealth in either direction. Indeed, the planner's problem is generically not convex, thus, we cannot contrast the planner's first-order conditions with the aggregate Euler equation (7) to assess the inefficiency as typically done in such models.

Returning to Figure 10, illustrates how the planner's solution and the competitive equilibrium differ during a payment-chain crisis, demonstrating the inefficient debt dynamics. Notice how in the extreme phase, the planner implements regressive policies increasing debt, which cuts back the chained expenditures of entrepreneurs—and increasing spot expenditures of creditors. Notice also how this regressive policy induces significant TFP gains, while barely affecting the entrepreneur's consumption, making it almost a strictly Pareto-improving policy. In the hysteresis phase, the wealth is redistributed toward entrepreneurs who now have room to make spot expenditures. Whereas TFP remains depressed for a while, TFP recovers substantially. The planner sacrifices a slight reduction in the creditor's consumption for a substantial increase in TFP. In the deleveraging phase, both solutions converge, indicating the productive and Pareto efficiency of both solutions in the long run.

5.4 Implication IV: Fiscal Policy (Bocola Effect)

In this section, we study fiscal policy and show that government spending has strikingly different effects depending on when it is paid and on the state of the economy. The key is that when governments pay upfront (spot), they inject immediate liquidity into payment chains. Chained government spending increases the chained share when it displaces private spot purchases; when it displaces private chained purchases, that share is unchanged. We call this the Bocola effect.[49]

To formalize this distinction, consider spot government expenditures G^s that require immediate payment, analogous to household spot purchases. The government can also make chained expenditures G^x contingent on receiving tax revenues, treating tax receipts as income flows. For simplicity, assume the government raises income taxes on entrepreneurs to finance expenditures under a balanced budget constraint. Government spending represents resource waste. The government enters as another agent with a budget constraint similar to that of entrepreneurs (with B_t=0). Chained expenditures now include those of the government:

$$\mu_t = q_t X_t^w + G_t^x.$$

[49] Economist Luigi Bocola conjectured this result would hold in our setting.

At fixed debt, when private chained purchases are positive, spot government spending reduces μ_t, whereas chained government spending replaces private chained expenditures one-for-one and leaves μ_t unchanged. When private purchases are all spot, chained government spending increases μ_t.

Proposition 5. (Government Multipliers): *Consider an infinitesimal rise in government expenditures evaluated at* $(G^x, G^s) = (0, 0)$. *Consumption and output multipliers are:* ***Efficient Region*** *(*$B \leq B^\star(\tilde{B})$*): Both expenditure types reduce welfare:*

$$\frac{\partial S}{\partial G^x} = \frac{\partial S}{\partial G^s} = -1, \quad \frac{\partial \mathcal{Y}}{\partial G^s} = 0, \quad \frac{\partial \mathcal{Y}}{\partial G^x} = \mathcal{A}(0^+) - 1 < 0.$$

Crisis Region *(*$B > B^\star(\tilde{B})$*): Spot and chained expenditures have opposite effects:*

$$\begin{aligned} \frac{\partial X}{\partial G^x} &= -\mathcal{A}(\mu), & \frac{\partial X}{\partial G^s} &= -\mathcal{A}(\mu)(1 + \epsilon_\mu^{\mathcal{A}}), \\ \frac{\partial \mathcal{Y}}{\partial G^x} &= 0, & \frac{\partial \mathcal{Y}}{\partial G^s} &= 1 - \mathcal{A}(\mu)(1 + \epsilon_\mu^{\mathcal{A}}) > 0. \end{aligned}$$

In words, the proposition states that in efficient regions, where no chained expenditures are made by the private sector, either form of government expenditures crowds out the entrepreneur's spot consumption. If the government expenditures are spot, they crowd out the entrepreneur's spot consumption. If the government spends through chained goods, it delays payments and reduces output, bearing all the losses directly.

During payment-chain crises in which the private sector makes chained expenditures, the effects differ. First, spot expenditures by the private sector do not change. Chained government expenditures remain harmful because they crowd out entrepreneur expenditures one-for-one. However, spot government expenditures create positive externalities by reducing average chain length. This shows a partial crowding in effect captured by $\epsilon_\mu^{\mathcal{A}}$, which is nonpositive. Moreover, when $\epsilon_\mu^{\mathcal{A}} < -1$, as in a deep crisis, it may be that spot government spending actually increases entrepreneur consumption despite higher taxes. The mechanism works through the speed-up of production, which more than offsets the cost of government resources.

This analysis yields a fresh perspective on fiscal stimulus. The conventional view is that a government that satisfies budget balance, as in this exercise, can increase output by stimulating aggregate demand in the presence of price rigidities and no tax distortions. Here, we stress that with payment delays, it is important that the government pays upfront for goods. A government that buys goods from firms that need funds to accelerate payments may inadvertently reduce output efficiency. During a payment-

chain crisis, paying upfront supports output by reducing the chained share. This dimension differs fundamentally from standard aggregate demand arguments—or Ricardian considerations. Instead, fiscal policy works through a payments channel: it stimulates output only when it has positive externalities on the private sector's payments. Consistent with this novel mechanism, Barrot and Nanda (2020) find evidence of a positive impact from faster government payments to small suppliers.

6 Conclusion

Financial crises have traditionally been understood through the lens of coordination failures and demand externalities. This paper identifies a distinct mechanism: economic activity can contract because payments must clear sequentially through a network before production and delivery can proceed. To our knowledge, this is the first paper to model bilateral payment obligations as a macroeconomic network and to test that mechanism using economy-wide data linking contractual terms, realized payment timing, delivery timing, prices, and firm outcomes. Together, the theory and evidence establish payment chains as a measurable channel through which financial disruptions can generate productivity losses.

In our theory, the core economic problem is that credit constraints lead to an inefficient distribution of funds, causing payment delays that disrupt production. Crises in this framework feature pecuniary externalities that arise endogenously from the payment structure. While our policy recommendations resemble those in environments with price rigidity, we underscore the importance of accelerating payments. For tractability, we make several key simplifications. We assume all transactions are bilateral and of identical size, and that delays occur within, rather than across, periods. We abstract from risk and inflation. Developing extensions with richer transaction sequences and heterogeneous delays is essential for quantitative applications. Nonetheless, we believe the core lessons about payment-chain disruptions are robust to such generalizations.

References

Aguiar, M., Amador, M., Farhi, E., Gopinath, G., 2015. Coordination and Crisis in Monetary Unions. The Quarterly Journal of Economics 130, 1727–1779. doi:10.1093/qje/qjv022.

Alvarez, F., Argente, D., Lippi, F., Méndez, E., Van Patten, D., 2023. Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments. Working Paper 31280. National Bureau of Economic Research. URL: http://www.nber.org/papers/w31280, doi:10.3386/w31280.

Alvarez, F., Barlevy, G., 2021. Mandatory disclosure and financial contagion. Journal of Economic Theory 194, 105237. doi:10.1016/J.JET.2021.105237.

Argente, D., Méndez, E., Van Patten, D., 2025. How New Products Enter Markets: Social Networks, Consumer Imports, and Retailer Learning .

Barrot, J.N., 2016. Trade Credit and Industry Dynamics: Evidence from Trucking Firms. The Journal of Finance 71, 1975–2016. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12371, doi:10.1111/jofi.12371.

Barrot, J.N., Nanda, R., 2020. The employment effects of faster payment: Evidence from the federal quickpay reform. The Journal of Finance 75, 3139–3173. doi:10.1111/jofi.12955.

Biais, B., Gollier, C., 1997. Trade credit and credit rationing. The Review of Financial Studies 10, 903–937.

Bigio, S., La'O, J., 2013. Financial frictions in production networks.

Bocola, L., 2022. Trade credit and financial amplification.

Boissay, F., Gropp, R., 2007. Trade credit defaults and liquidity provision by firms. European Central Bank working paper series, 753.

Brunnermeier, M.K., Sannikov, Y., 2017. The I-Theory of Money.

Buera, F.J., Shin, Y., 2013. Financial frictions and the persistence of history: A quantitative exploration. Journal of Political Economy 121, 221–272. doi:10.1086/670271.

Burkart, M., Ellingsen, T., 2004. In-kind finance: A theory of trade credit. American Economic Review 94, 569–590.

Cerdas, A.M., Melegatti, C., 2014. El sistema financiero costarricense en los últimos 25 años: sistema de pagos. Number 2 in Monografía, Academia de Centroamérica, San José, Costa Rica. 28 cm.

Cole, H.L., Kehoe, T.J., 2000. Self-Fulfilling Debt Crises. The Review of Economic Studies 67, 91–116. URL: http://restud.oxfordjournals.org/content/67/1/91.abstract, doi:10.1111/1467-937X.00123.

Costello, A.M., 2020. Credit Market Disruptions and Liquidity Spillover Effects in the Supply Chain. Journal of Political Economy 128, 3434–3468.

Cuñat, V., 2007. Trade credit: Suppliers as debt collectors and insurance providers. Review of Financial Studies 20, 491–527.

Diamond, P.A., 1982. Aggregate Demand Management in Search Equilibrium. Journal of Political Economy 90, 881–894.

Eggertsson, G.B., Krugman, P., 2012. Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. The Quarterly Journal of Economics 127, 1469–1513. URL: https://doi.org/10.1093/qje/qjs023, doi:10.1093/qje/qjs023. _eprint: https://academic.oup.com/qje/article-pdf/127/3/1469/30457060/qjs023.pdf.

Elliott, M., Golub, B., Jackson, M.O., 2014. Financial networks and contagion. American Economic Review 104, 3115–53. doi:10.1257/AER.104.10.3115.

Elliott, M., Golub, B., Leduc, M.V., 2022. Supply network formation and fragility. American Economic Review 112, 2701–47. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20210220, doi:10.1257/aer.20210220.

Freeman, S., 1996. The Payments System, Liquidity, and Rediscounting. The American Economic Review 86, 1126–1138. URL: http://www.jstor.org/stable/2118282. publisher: American Economic Association.

Green, E.J., 1999. Money and Debt in the Structure of Payments. Federal Reserve Bank

of Minneapolis Quarterly Review 23, 13–29.
Guerrieri, V., Lorenzoni, G., 2009. Liquidity and Trading Dynamics. Econometrica 77. URL: http://doi.wiley.com/10.3982/ECTA7231, doi:10.3982/ECTA7231.
Guerrieri, V., Lorenzoni, G., 2017. Credit Crises, Precautionary Savings, and the Liquidity Trap*. The Quarterly Journal of Economics 132, 1427–1467. URL: https://doi.org/10.1093/qje/qjx005, doi:10.1093/qje/qjx005.
Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I., 2022. Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? AER .
Hardy, B., Simonovska, I., Saffie, F.E., 2022. Economic stabilizers in emerging markets:the case for trade credit.
Jacobsen, Tor; von Schedvin, E., 2015. Trade credit and the propagation of corporate failure: An empirical analysis. Econometrica 83, 1315–1371.
Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. American economic review 95, 161–182.
Jordà, Ò., Taylor, A.M., 2025. Local projections. Journal of Economic Literature 63, 59–110.
Kiyotaki, N., Moore, J., 1997a. Credit Chains. Publisher: London School of Economics.
Kiyotaki, N., Moore, J., 1997b. Credit Cycles. Journal of Political Economy 105, 211–248. URL: http://www.jstor.org/stable/2138839. publisher: The University of Chicago Press.
Kiyotaki, N., Wright, R., 1989. On Money as a Medium of Exchange. Journal of Political Economy 97, 927–954. Publisher: The University of Chicago Press.
Kopytov, A., Mishra, B., Nimark, K., Taschereau-Dumouchel, M., 2022. Endogenous production networks under supply chain uncertainty.
Lagos, R., Rocheteau, G., 2009. Liquidity in Asset Markets with Search Frictions. Econometrica 77, 403–426.
Lagos, R., Rocheteau, G., Weill, P.O., 2011. Crises and Liquidity in OTC Markets. Journal of Economic Theory 146, 2169–2205.
Lagos, R., Wright, R., 2005. Unified Framework for Monetary Theory and Policy Analysis. Journal of Political Economy 113, 463–484.
La'O, J., 2015. A traffic jam theory of recessions.
Levinsohn, J., Petrin, A., 2003. Estimating production functions using inputs to control for unobservables. The Review of Economic Studies 70, 317–341. doi:10.1111/1467-937X.00246.
Li, Y., Rocheteau, G., Weill, P.O., 2012. Liquidity and the Threat of Fraudulent Assets. Journal of Political Economy 120, 815–846.
Lippi, F., Ragni, S., Trachter, N., 2015. Optimal monetary policy with heterogeneous money holdings. Journal of Economic Theory 159, 339 – 368. URL: http://www.sciencedirect.com/science/article/pii/S0022053115001283, doi:https://doi.org/10.1016/j.jet.2015.07.005.
Lucas, R.E., Stokey, N.L., 1987. Money and interest in a cash-in-advance economy. Econometrica 55, 491–513. URL: http://www.jstor.org/stable/1913597. publisher: The Econometric Society.
Meza, F., Quintin, E., 2007. Factor utilization and the real impact of financial crises. The B.E. Journal of Macroeconomics 7. URL: https://doi.org/10.2202/1935-1690.1593, doi:doi:10.2202/1935-1690.1593.

Méndez, E., Van Patten, D., 2025. Voting on a trade agreement: Firm networks and attitudes towards openness. The Review of Economic Studies , rdaf005URL: https://doi.org/10.1093/restud/rdaf005, doi:10.1093/restud/rdaf005.

Nosal, E., Rocheteau, G., 2011. Money, Payment, and Liquidity. MIT Press.

Oberfield, E., 2013. Productivity and misallocation during a crisis: Evidence from the chilean crisis of 1982. Review of Economic Dynamics 16, 100–119. URL: https://www.sciencedirect.com/science/article/pii/S1094202512000634, doi:https://doi.org/10.1016/j.red.2012.10.005.

Oberfield, E., 2018. A theory of input-output architecture. Econometrica 86, 559–589. URL: https://onlinelibrary.wiley.com/doi/10.3982/ECTA10731, doi:10.3982/ECTA10731.

Petersen, M.A., Rajan, R.G., 1997. Trade credit: Theories and evidence. Review of Financial Studies 10, 661–691.

Pratap, S., Urrutia, C., 2012. Financial frictions and total factor productivity: Accounting for the real effects of financial crises. Review of Economic Dynamics 15, 336–358. URL: https://www.sciencedirect.com/science/article/pii/S1094202511000408, doi:https://doi.org/10.1016/j.red.2011.09.003.

Rocheteau, G., 2011. Payment and Liquidity under Adverse Selection. Journal of Monetary Economics 58.

Rocheteau, G., Wong, T.N., Weill, P.O., 2016. Working through the Distribution: Money in the Short and Long Run.

Shi, S., 1997. A divisible search model of fiat money. Econometrica 65, 75–102.

Taschereau-Dumouchel, M., 2022. Cascades and fluctuations in an economy with an endogenous production network.

Townsend, R.M., 1980. Models of monetary economies , 265–303.

Woodford, M., 2022. Effective demand failures and the limits of monetary stabilization policy. American Economic Review 112, 1475–1521. doi:10.1257/AER.20201529.

Appendix (Not intended for publication)

Table of Contents

Appendix to Payment-Chain Block (Section 2) **1**

A.1 Microfoundations for formula for y_n 1

A.2 Assignment Functions and Outcomes (Proof of Proposition 1) 4

A.3 Related Results Used Elsewhere 14

Extensions of the Payment-Chain Block **16**

B.1 Correlated Assignments 18

B.2 Pooled Chained Orders 21

B.3 Endogenous Reallocation 23

B.4 Default and Cascading 25

B.5 Partial Delays 26

Appendix to the Empirical Analysis (Section 3) **29**

C.1 Validation: Incoming liquidity predicts outgoing payments 29

C.2 Delivery in Retail 37

C.3 Validating Shipping-Payment Dates with Retail Product Entry 37

C.4 Price Premium in the Shipping-Observed Sample 38

C.5 Timing-Based Chained Payments and Prices 40

C.6 Firm-Level TFP Estimation 40

C.7 Earlier-Outstanding-Invoice Placebo 41

C.8 Complementary Timing-Based Evidence from the Desyfin Freeze . . . 43

Proofs for the Partial Equilibrium Analysis (Section 4) **45**

D.1 Entrepreneur's Problem: Details 45

D.2 Proof of Proposition 2 (Entrepreneur's Euler Equation) 47

D.3 Proof of Proposition 3 (Entrepreneur's Stationary Solution) 50

Proofs in the General Equilibrium Analysis (Section 5) **55**

E.1 Proof of Proposition 4 55

E.2 Analysis: Preliminary Observations 56

E.3 Derivation of the map $\mathcal{B}$ 58
E.4 Proof of Corollary 1 62
E.5 Proof of Corollary 2 62

Constrained Efficiency Analysis (Section 5) **63**

F.1 Characterization of the Primal Problem 63
F.2 Proof of Proposition 8 (Solution of the Primal Problem) 65
F.3 Discussion of Planner's Solution. 74
F.4 Decentralization with Taxes 77
F.5 Proof Ramsey-Primal Equivalence - Proposition 9 79
F.6 Derivation of the Bocola Effect - Proof of Proposition 5 82

A Appendix to Payment-Chain Block (Section 2)

A.1 Microfoundations for formula for y_n

A.1.1 Value Added - Delayed Production

The goal of this section is to provide a setting where value added for an order in position n is given by $y_n = \delta^n$ and the transaction time is $\tau_n = 1 - \delta^n$. This formulation ensures that all transactions occur within the unit time interval $[0, 1]$, with $\tau_n \to 1$ as $n \to \infty$. This within-period structure allows us to embed payment chains into an infinite-horizon discrete-time framework, as we do in Section 5. Specifically, the entire payment-chain network—from the first spot order at $\tau_0 = 0$ to the limiting transaction as $n \to \infty$—forms and settles within each integer period. This approach preserves analytical tractability while capturing the essential dynamics of sequential payments and production delays.

Time and Production. Value added is generated by translating inputs into final goods. To produce final goods, input endowments must be transferred to the agent placing the order for the good. Recall that there are N orders in the production network, indexed by their position $k = 1, 2, \ldots, N$ in the bilateral production network, induced by an assignment rule $\mathcal{P}$, as in the main text. Within this production network, payment chains emerge. Further recall that we index orders by their position $n = 0, 1, 2, \ldots$ within a payment chain. An order at production position k corresponds to some position n in its payment chain.

Value added. As in the main text, orders are for customized intermediate inputs: An order at position k in the bilateral production network is supplied by an order in position $k + 1$ and so on. The customized nature of these goods means they have no resale value. Moreover, recall that the production network and the income-expenditure identity induce a network of payment chains.

An order in position n in a payment chain is fulfilled at a time $\tau_n \in [0, 1]$. This means that the intermediate input is available to the agent placing the order at τ_n. The technology to produce final goods (gross output) is given by:

$$y_n = \int_{\tau_n}^{1} \min\{x_t, h_t\} dt. \tag{10}$$

Here, h_t stands for the labor input per unit of time that can be used by the agent placing

the order to produce final goods. We assume that the agent placing the order has an endowment of labor $h_t \leq 1$ to process each order.[50] In turn, x_t represents the use of the intermediate input per unit of time. The total input endowment is 1 unit, and input use must satisfy:

$$\int_{\tau_n}^{1} x_t dt \leq 1. \tag{11}$$

Because once payments for orders have been executed and there's no resale use for the inputs, the only decision concerns the use of h_t and x_t. To maximize value added y_n, requires setting $h_t = x_t = 1$. Thus, output is given by $y_n = \int_{\tau_n}^{1} \min\{x_t, h_t\} dt = 1 - \tau_n$. Suppose that transaction times indeed occur at times $\tau_n = 1 - \delta^n$. Thus, in that case, $y_n = \delta^n$.

These calculations imply that for an order in position n, the share $\tau_n = 1 - \delta^n$ of its input is wasted.[51] Next, we derive an explanation of why the transaction times are indeed instances $\tau_n = 1 - \delta^n$.

Delays arising from fraud prevention. We consider a scenario where delays in transferring funds and, therefore, delays in transferring inputs are the result of a payment system designed to deter fraud.

To that end, we consider the possibility of fraudulent orders. On the one hand, we have *buyer-side fraud*: a population of fraudster agents can place fraudulent chain orders. These fraudulent chained orders falsely claim to have incoming payments, but in reality, they have nothing to sell and know that no incoming funds will arrive. The presence of these agents implies that to prevent a flooding of buyer-side fraud, any transfer of inputs must require *proof of funds*; otherwise, trade would break down.

On the other hand, we consider the possibility of *seller-side fraud*: a population of fraudster agents that can provide fraudulent inputs. These inputs generate no value added. However, by selling these inputs, fraudster sellers can obtain funds to buy inputs while harming their buyers. Trade is possible only if goods can be inspected, which takes time.

To enable trade despite these fraudulent orders and suppliers, transactions are conducted through a payment system that enables an escrow-inspection mechanism:

1. Once an order at position n receives funds at τ_n, the buyer can deposit them into an escrow mechanism. The funds can be proven to exist, but will not be released

[50] We could alternatively model a competitive market for labor with elastic supply, which would not change the main results leading to the same outcome.

[51] In the extensions section, we discuss variations so that the unused inventory is not waste.

until the buyer sends out a payment instruction.

2. Once the funds are in escrow, the seller (at position $k+1$ in the production network) verifies that the buyer is not fraudulent. The seller delivers the inventory immediately. As a result, the buyer has the inventory at time τ_n. The buyer begins production immediately. The presence of verifiable funds deters buyer-side fraud.

3. The buyer requires inspection to deter seller-side fraud. To do so, the buyer must produce a fraction $(1-\delta)$ of output as a quality verification sample. This inspection requires to first observe the share $(1-\delta)$ of the output that will be produced.

4. The escrowed funds are released to the seller at position $n+1$ in the payment chain. If the seller proves the inputs are fraudulent, the funds are returned to the buyer. Whenever the full input units arrive in some $\tau > 0$ not all the input can be used. In particular, the buyer will use only share $x = 1-\tau_n$ of the input. Thus, the inspection requires $(1-\delta)(1-\tau_n)$ inspection time given that the output is just $(1-\tau_n)$.

Notice that the parameter $\delta \in (0,1)$ represents inspection efficiency or verification technology. Furthermore, observe that if buyers are non-fraudulent, they execute payment. We are implicitly assuming commitment on the side of honest buyers: they pay the seller even when having the inventory and despite how late they may be in the chain.

Derivation of Timing of Payments. Consider position n in the chain: the buyer receives inventory at time τ_n. Since value added is: $y_n = 1-\tau_n$ of which $(1-\delta)(1-\tau_n)$ is the inspection time, payment is released to position $n+1$ at: $\tau_n + (1-\delta)y_n$. Thus, we obtain the recursion:

$$\tau_{n+1} = \tau_n + (1-\delta)y_n = \tau_n + (1-\delta)(1-\tau_n) = (1-\delta) + \delta\tau_n \tag{12}$$

The initial condition is $\tau_0 = 0$ (spot orders begin immediately with external funding). Thus, we obtain the recursion:

$$\begin{aligned}
\tau_1 &= 1-\delta \\
\tau_2 &= (1-\delta) + \delta(1-\delta) = 1-\delta^2 \\
\tau_3 &= (1-\delta) + \delta(1-\delta^2) = 1-\delta^3.
\end{aligned}$$

Hence, we establish the desired result that:

$$y_n = 1 - \tau_n = \delta^n. \tag{13}$$

Examples of escrow-inspection mechanisms in practice. Our setting is admittedly simplistic. However, the payment delays that occur through escrow-inspection mechanisms we model—involving proof of funds, escrow, and quality verification—are ubiquitous in real-world payment systems.

Contracts that involve shipping can illustrate these requirements when the buyer commits funds to escrow before delivery and the supplier is paid only after inspection and acceptance. The timing in our model further requires that inspection occur during production after delivery.

Traditional payment technologies, such as checks, embody a similar logic. When a buyer pays by check, processing delays provide a window for inspection before funds transfer. If goods prove defective, the buyer can stop payment; if the seller lacks sufficient backing, the check bounces, revealing fraud. The check-clearing delay thus naturally accommodates inspection before irrevocable settlement.

These mechanisms are equally central to digital commerce. Online marketplaces such as eBay and MercadoLibre operate buyer and seller protection programs that hold payments in escrow until buyers confirm receipt and quality. In business-to-business transactions, vendors commonly require bank confirmation letters proving funds exist before shipping goods, while buyers release final payment only after inspection and acceptance. Blockchain technologies and digital wallets extend these principles further: wallet balances are publicly verifiable, and customers can demonstrate ownership and proof of funds by initiating small test transactions, providing transparent verification to all parties in the chain.

Under the assumed inspection and escrow restrictions, fraud prevention entails delays in payment execution. While we do not study optimal payment system design in this paper—an interesting topic in its own right—our framework captures a tension between transaction speed and fraud prevention.

A.2 Assignment Functions and Outcomes (Proof of Proposition 1)

Proposition 1 presents several results for random and efficient assignments. We start with the formulas for random assignments

A.2.1 Random Assignments.

Chain-Length Distribution. We first show that the distribution of payment-chain lengths is geometric.

Lemma 2. *For the random assignment, the chain-length distribution is geometric.*

Proof. Recall that funded nodes are placed randomly along the production circle, with $N^x/N \to \mu$. For any fixed number of successive positions, sampling without replacement converges to independent sampling as $N \to \infty$. Thus, the probability that the next node is chained converges to μ, and the probability that it is funded converges to $1-\mu$.

By construction, every chained order must be funded through exactly one production link, creating a directed path that ultimately terminates at a spot order. This structure ensures that every node belongs to exactly one payment chain, and each chain originates from a unique spot order. Therefore, we have exactly N^s payment chains, which we index by $i \in \mathcal{N}^s$.

Because there is a one-to-one map from spot orders to chains, we use an index for all possible chains by defining chain i as the payment chain that starts with spot order i. For chain i, corresponding to an initial spot order i, we define its length as the number of chained orders it contains (equivalently, the number of edges in the directed path). We now derive the distribution of chain lengths.

Consider chain i: It is of **length 0** if spot order i orders from a node associated with a spot order, which occurs with probability $1-\mu$. It is of **length 1** the first order in the chain orders from a chained order, which happens with probability μ, whereas that order is not linked, which occurs with probability $1-\mu$. Thus, the overall probability is $\mu(1-\mu)$. Proceeding by induction, for a chain of **length** n, extending this reasoning, there must be n consecutive chained orders, followed by a termination (no further links). By independence, this occurs with probability $\mu^n(1-\mu)$. Therefore, if L_i denotes the length of chain i, we have:

$$\mathbb{P}(L_i = n) = \mu^n(1-\mu) \quad \text{for } n \in \{0,1,2,\ldots\},$$

i.e., the geometric distribution with parameter $1-\mu$ (using the convention where support begins at 0). □

Output and TFP. Next, we derive the distribution of positions for the geometric distribution.

Lemma 3. *When the chain length distribution is geometric, the distribution of positions is also geometric.*

Proof. For the geometric distribution:

$$G(n-1) = (1-\mu)\sum_{m=0}^{n-1}\mu^n = (1-\mu^n), \quad n \geq 1.$$

Therefore, $1 - G(n-1) = \mu^n$, and, thus,

$$\sum_{n}^{\infty}(1 - G(n-1)) = \sum_{n=0}^{\infty}\mu^n = \frac{1}{1-\mu}.$$

Thus, we have that:

$$h(n) = \frac{\mu^n}{1/(1-\mu)} = (1-\mu)\mu^n.$$

Thus, the distribution of positions follows a geometric distribution with the same parameter $1-\mu$, completing the proof. □

Next, we derive the formula for TFP, $\mathcal{A}$, and output $\mathcal{Y}$ using two approaches. We first do so by employing the distribution of positions.

Lemma 4. *For the random assignment we have that* $\mathcal{Y} = \frac{1-\mu}{1-\delta\mu}$ *and* $\mathcal{A} = \delta\frac{1-\mu}{1-\delta\mu}$.

Calculations using the Position Distribution h.

Proof. Given that the distribution of positions $h(n) = (1-\mu)\mu^n$ the average quantity of final goods,

$$\mathcal{Y} \equiv \mathbb{E}[y_n] = \mathbb{E}[\delta^n] = \sum_{n=0}^{\infty}(1-\mu)\mu^n\delta^n = \frac{1-\mu}{1-\delta\mu}.$$

Then, it follows that $\mathcal{Y} = (1-\mu) + \mu\mathcal{A} \rightarrow \mathcal{A} = \delta\frac{1-\mu}{1-\delta\mu}$. □

Calculation using the Length Distribution.

Alternative Proof via Weighted Averaging. We demonstrate a general approach for deriving position distributions from chain length distributions, applicable even when direct computation is complex. We illustrate with the geometric case where $\mathbb{P}(L = n) = (1-\mu)\mu^n$.

Step 1: Expected production per chain. For a chain of length $n \geq 1$, the cumulative production discount is:

$$E[y_n|\text{length} = n] = \sum_{i=1}^{n} \delta^i = \frac{\delta(1-\delta^n)}{1-\delta}$$

where $\delta \in (0,1)$ represents the production efficiency at each stage.

Step 2: Average chain length. The expected number of chained orders per chain (excluding spot orders) is:

$$\bar{n}_x = \sum_{n=1}^{\infty} n \cdot \mathbb{P}(L=n) = \sum_{n=1}^{\infty} n(1-\mu)\mu^n \tag{14}$$

$$= (1-\mu)\mu \sum_{n=1}^{\infty} n\mu^{n-1} \tag{15}$$

$$= (1-\mu)\mu \cdot \frac{d}{d\mu}\left(\frac{1}{1-\mu}\right) \tag{16}$$

$$= (1-\mu)\mu \cdot \frac{1}{(1-\mu)^2} = \frac{\mu}{1-\mu} \tag{17}$$

Step 3: Position distribution via weighted averaging. To find the distribution of positions, we weight each chain length by its contribution to the total mass of chained orders:

$$\text{Weight for length } n = \frac{n \cdot \mathbb{P}(L=n)}{\bar{n}_x}$$

The probability that a randomly selected chained order is at position k equals:

$$\mathbb{P}(X=k|X>0) = \sum_{n=k}^{\infty} \frac{(1-\mu)\mu^n \cdot 1}{\mu/(1-\mu)} = (1-\mu)^2 \sum_{n=k}^{\infty} \mu^{n-1} = (1-\mu)\mu^{k-1}$$

Accounting for spot orders (position 0), we have:

$$\mathbb{P}(X=k) = \begin{cases} 1-\mu & \text{if } k=0 \\ \mu \cdot (1-\mu)\mu^{k-1} = (1-\mu)\mu^k & \text{if } k \geq 1 \end{cases}$$

This confirms that positions follow a geometric distribution with parameter $1-\mu$.

□

Part 2. Limits. We first consider the limit as $\mu \to 0$:

$$\lim_{\mu \to 0} \mathcal{A}(\mu; \delta) = \delta \lim_{\mu \to 0} \left(\frac{1-\mu}{1-\delta\mu} \right) = \delta.$$

For output:

$$\lim_{\mu \to 0} \mathcal{Y}(\mu; \delta) = \lim_{\mu \to 0} (1-\mu) + \mu \mathcal{A}(\mu; \delta) = 1.$$

Next, we consider the limit as $\mu \to 1$:

$$\lim_{\mu \to 1} \mathcal{A}(\mu; \delta) = \delta \lim_{\mu \to 1} \left(\frac{1-\mu}{1-\delta\mu} \right) = 0.$$

For output:

$$\lim_{\mu \to 1} \mathcal{Y}(\mu) = \lim_{\mu \to 1} (1-\mu) + \lim_{\mu \to 1} \mu \lim_{\mu \to 1} \mathcal{A}(\mu) = 0.$$

Next, we consider the limit as $\delta \to 0$:

$$\lim_{\delta \to 0} \mathcal{A}(\mu; \delta) = \lim_{\delta \to 0} \frac{\delta - \delta\mu}{1-\mu\delta} = 0.$$

For output,

$$\lim_{\delta \to 0} \mathcal{Y}(\mu) = (1-\mu) + \mu \lim_{\delta \to 0} \mathcal{A}(\mu; \delta) = (1-\mu).$$

Finally, we consider the limit as $\delta \to 1$:

$$\lim_{\delta \to 1} \mathcal{A}(\mu; \delta) = \lim_{\delta \to 1} \delta \cdot \lim_{\delta \to 1} \frac{1-\mu}{1-\delta\mu} = 1.$$

For output,

$$\lim_{\delta \to 1} \mathcal{Y}(\mu) = (1-\mu) + \mu \lim_{\delta \to 1} \mathcal{A}(\mu; \delta) = 1.$$

Part 3. Monotonicity and concavity. We establish the following result:

Lemma 5. *$\mathcal{Y}(\mu), \mathcal{A}(\mu)$ are strictly decreasing and concave in μ.*

Proof. Next, we investigate the monotonicity of $\mathcal{A}$ and $\mathcal{Y}$. We do so by signing the derivatives.

$$\mathcal{Y}(\mu) = (1-\mu)\left(1 + \frac{\delta\mu}{1-\delta\mu}\right) = \frac{1-\mu}{1-\delta\mu}.$$

Therefore, we obtain,

$$\begin{aligned}\mathcal{Y}_\mu &= \mathcal{Y}(\mu)\left(\frac{\delta}{1-\delta\mu}-\frac{1}{1-\mu}\right)\\ &= -\frac{1-\delta}{(1-\delta\mu)^2}<0.\end{aligned}$$

We also have that, $\mathcal{Y}_\mu = -(1-\mathcal{A}(\mu))+\mu\mathcal{A}_\mu$. Then,

$$1-\mathcal{A}(\mu)=\frac{1-\delta\mu-\delta+\delta\mu}{1-\delta\mu}=\frac{1-\delta}{1-\delta\mu}.$$

Therefore:

$$\mu\mathcal{A}_\mu=\left(\frac{1-\delta}{1-\delta\mu}-\frac{1-\delta}{(1-\delta\mu)^2}\right)=-\frac{1-\delta}{1-\delta\mu}\left(\frac{\delta\mu}{1-\delta\mu}\right)<0.$$

Thus:

$$\mathcal{A}_\mu=-\frac{1-\delta}{1-\delta\mu}\left(\frac{\delta}{1-\delta\mu}\right). \tag{18}$$

Next, we perform the convexity analysis by computing second derivatives. Using (19) we have that:

$$\mathcal{Y}_{\mu\mu} = -2\delta\frac{1-\delta}{(1-\delta\mu)^3}<0.$$

Thus, $\mathcal{Y}(\mu)$ is strictly concave.

From $\mathcal{Y}_{\mu\mu}=2\mathcal{A}_\mu+\mu\mathcal{A}_{\mu\mu}$, we obtain

$$\begin{aligned}\mu\mathcal{A}_{\mu\mu} &= \mathcal{Y}_{\mu\mu}-2\mathcal{A}_\mu\\ &= -2\delta\frac{1-\delta}{(1-\delta\mu)^3}+2\delta\frac{1-\delta}{(1-\delta\mu)^2}\\ &= -2\delta\frac{1-\delta}{(1-\delta\mu)^3}\delta\mu \quad\Rightarrow\quad \mathcal{A}_{\mu\mu}=-2\delta^2\frac{1-\delta}{(1-\delta\mu)^3}<0.\end{aligned}$$

Thus, we establish that both functions are strictly monotonically decreasing and concave in μ.

□

We also have monotonicity in δ, but a different convexity property.

Lemma 6. *$\mathcal{Y}(\mu),\mathcal{A}(\mu)$ are strictly increasing in δ. In turn, both $\mathcal{Y}(\mu)$ and $\mathcal{A}(\mu)$ are strictly convex in δ.*

Proof. Note that:

$$\mathcal{Y}_\delta \;=\; \mathcal{Y}(\mu)\left(\frac{\mu}{1-\delta\mu}\right) > 0.$$

Moreover, $\mu\mathcal{A}_\delta = \mathcal{Y}_\delta \to \mathcal{A}_\delta = \frac{\mathcal{Y}(\mu)}{1-\delta\mu} > 0$.

As for the convexity analysis:

$$\mathcal{Y}_{\delta\delta} \;=\; \mathcal{Y}_\delta\left(\frac{\mu}{1-\delta\mu}\right) + \mathcal{Y}(\mu)\frac{\mu^2}{(1-\delta\mu)^2} = 2\mathcal{Y}(\mu)\frac{\mu^2}{(1-\delta\mu)^2} > 0.$$

From here, we have that, $\mathcal{A}_{\delta\delta} = \mathcal{Y}_\delta\left(\frac{1}{1-\delta\mu}\right) + \mathcal{Y}(\mu)\frac{\mu}{(1-\delta\mu)^2} = 2\frac{\mu(1-\mu)}{(1-\delta\mu)^3} > 0$. This establishes the strict monotonicity and convexity properties. □

A.2.2 Efficient Assignment

We now prove the formulas for the chain-size distribution for the efficient assignment.

Lemma 7. *Let,* $n = \left\lfloor \frac{\mu}{1-\mu} \right\rfloor$ *and* $p = n + 1 - \frac{\mu}{1-\mu}$. *Then, under the efficient assignment* $g(n) = p, \quad g(n+1) = 1-p$.

Proof. Consider the problem of allocating $N^x = \mu N$ chained orders across $N^s = (1-\mu)N$ chains to maximize total output.

Step 1: Marginal productivity and reallocation.

Consider any assignment with two chains of lengths n and n' where $n' > n+1$. We can strictly increase output by moving the order from the final position in the chain of length n' to adding a link to the chain of length n. The final goods obtain by that order is:

$$\delta^{n+1} - \delta^{n'} = \delta^{n+1}(1 - \delta^{n'-n-1}) > 0.$$

The final goods of no other order is not affected.

Step 2: Characterization of efficient allocation. By Step 1, in any efficient allocation, chain lengths can differ by at most 1. Therefore, the distribution must be:

$$g(k) = \begin{cases} p & \text{if } k = n \\ 1-p & \text{if } k = n+1 \\ 0 & \text{otherwise} \end{cases}$$

for some integer $n \geq 0$ and $p \in [0,1]$. Otherwise, it's possible to move a mass of orders and increase the final output.

Step 3: Determining n and p. The total number of chained orders must equal N^x:

$$N^s[pn + (1-p)(n+1)] = N^x.$$

Dividing both sides by N^s:

$$pn + (1-p)(n+1) = \frac{N^x}{N^s} = \frac{\mu N}{(1-\mu)N} = \frac{\mu}{1-\mu}.$$

Simplifying:

$$n + (1-p) = \frac{\mu}{1-\mu}.$$

Therefore:

$$p = n + 1 - \frac{\mu}{1-\mu}.$$

Step 4: Integer constraints. For $p \in [0,1]$, we need:

$$0 \leq n + 1 - \frac{\mu}{1-\mu} \leq 1. \tag{19}$$

We choose the lower integer n so that:

$$n \leq \frac{\mu}{1-\mu} < n + 1.$$

Hence, $n = \left\lfloor \frac{\mu}{1-\mu} \right\rfloor$ and $p = n + 1 - \frac{\mu}{1-\mu}$. When $\mu/(1-\mu)$ is an integer, $p = 1$ and all chains have length n.

Step 5: Final distribution. The efficient allocation has:

- A fraction $p = \left\lfloor \frac{\mu}{1-\mu} \right\rfloor + 1 - \frac{\mu}{1-\mu}$ chains with length $n = \left\lfloor \frac{\mu}{1-\mu} \right\rfloor$
- A fraction $1 - p = \frac{\mu}{1-\mu} - \left\lfloor \frac{\mu}{1-\mu} \right\rfloor$ of chains with length $n + 1$

Note that for $\mu < \frac{1}{2}$, the mass of chains of length 1 is $1 - p = \frac{\mu}{1-\mu} = \frac{N^x}{N^s}$, precisely the ratio of chained orders to spot orders.

This completes the characterization of the efficient chain-length distribution. □

The unconditional distribution of positions $H(k)$ is thus given by:

$$H(k) = \begin{cases} \frac{1}{1+n+(1-p)} & \text{if } 0 \leq k \leq n \\ \frac{1-p}{1+n+(1-p)} & \text{if } k = n+1 \\ 0 & \text{otherwise} \end{cases}$$

where $n = \left\lfloor \frac{\mu}{1-\mu} \right\rfloor$. Next, we derive the formula for $\mathcal{A}$ under the efficient assignment.

Lemma 8. *Under the efficient assignment:*

$$\mathcal{A}(\mu;\delta) = \delta \left(p\frac{1-\delta^n}{1-\delta} + (1-p)\frac{1-\delta^{n+1}}{1-\delta} \right) \frac{1-\mu}{\mu}.$$

Furthermore:

$$\mathcal{Y}(\mu;\delta) = (1-\mu)\left(1 + \delta\left(p\frac{1-\delta^n}{1-\delta} + (1-p)\frac{1-\delta^{n+1}}{1-\delta} \right)\right).$$

Proof. Note that in a chain of length n, the average output among chained orders: $\frac{1}{n}\sum_{i=1}^{n}\delta^i = \frac{\delta}{n}\frac{1-\delta^n}{1-\delta}$. The mass of chained orders is: $pn + (1-p)(n+1) = \frac{\mu}{1-\mu}$. From here, we compute output per chained order:

$$\mathcal{A}(\mu) = \delta \left(p\frac{1-\delta^n}{1-\delta} + (1-p)\frac{1-\delta^{n+1}}{1-\delta} \right) / \left(\frac{\mu}{1-\mu} \right), \tag{20}$$

yielding the formula. □

Monotonicity and concavity. Next, we prove the established monotonicity and continuity.

Lemma 9. *$\mathcal{A}(\mu)$ and $\mathcal{Y}(\mu)$ are continuous and monotone decreasing under the efficient allocation. Moreover, $\mathcal{Y}(\mu)$ is concave in μ: it is linear on each interval where $n = \lfloor \mu/(1-\mu) \rfloor$ is constant, and its slope falls at each kink. By contrast, $\mathcal{A}(\mu)$ is strictly convex in μ on each such interval with $n \geq 1$, so $\mathcal{A}$ is neither globally concave nor globally convex.*

Proof. Recall that under efficient allocation, average output per chained order is:

$$\mathcal{A}(\mu) = \frac{\delta}{1-\delta}\left(p(1-\delta^n) + (1-p)(1-\delta^{n+1}) \right)\frac{1-\mu}{\mu}$$

where $n = \left\lfloor \frac{\mu}{1-\mu} \right\rfloor$ and $p = n + 1 - \frac{\mu}{1-\mu}$.

Step 1: Continuity at integer values. When $\frac{\mu}{1-\mu} = k$ for integer k, all chains have length k, giving:

$$\mathcal{A}(\mu) = \frac{\delta(1-\delta^k)}{1-\delta} \cdot \frac{1-\mu}{\mu}$$

As μ approaches this integer point from below, $n = k-1$ and $p \to 0$, yielding the same limit. Thus $\mathcal{A}(\mu)$ is continuous everywhere.

Step 2: Derivative between integer points. For $\frac{\mu}{1-\mu} \notin \mathbb{Z}$, the floor function is locally constant, so:

$$\frac{\partial \mathcal{A}}{\partial \mu} = \frac{\delta}{1-\delta} \cdot \frac{\partial}{\partial \mu}\left(\frac{1-\mu}{\mu}\right) \cdot [(1-\delta^n) - n(1-\delta)\delta^n]$$

Since $\frac{\partial}{\partial \mu}\left(\frac{1-\mu}{\mu}\right) = -\frac{1}{\mu^2} < 0$, it suffices to show that the bracket is nonnegative:

$$(1-\delta^n) - n(1-\delta)\delta^n \geq 0$$

Step 3: Verifying the inequality. Rearranging:

$$(1-\delta^n) - n(1-\delta)\delta^n = 1 - (1+n)\delta^n + n\delta^{n+1} \geq 0.$$

For $n = 0$, the expression is identically zero. For $n \geq 1$, define $f(\delta) = 1 + n\delta^{n+1} - (1+n)\delta^n$. We have $f(0) = 1$, $f(1) = 0$, and $f'(\delta) = n(n+1)\delta^{n-1}(\delta - 1) < 0$ for $\delta \in (0,1)$. Thus, $f(\delta) > 0$ for $n \geq 1$, establishing strict decrease of $\mathcal{A}$ between kinks when $\mu > 1/2$.

Step 4: Global monotonicity. Although the derivative has discontinuities at integer values of $\frac{\mu}{1-\mu}$, the function is continuous and weakly decreasing between these points. At integer points, the left and right limits coincide, ensuring global monotonicity.

Therefore, $\mathcal{A}(\mu)$ is weakly decreasing on $(0,1)$: it equals δ for $0 < \mu \leq 1/2$ and is strictly decreasing for $1/2 < \mu < 1$.

Step 4: Monotonicity in δ. This follows immediately since δ^a is increasing in δ for any integer a.

Step 5: Curvature in μ. Fix n and let $r \equiv \frac{\mu}{1-\mu} \in (n, n+1)$, so that $p = n + 1 - r$ varies with μ. Let $S_n \equiv \sum_{k=1}^{n} \delta^k$. Output per chain is

$$pS_n + (1-p)S_{n+1} = S_n + (r-n)\delta^{n+1},$$

which is linear in r. Dividing by the number of chained orders per chain, r, and using $1/r = (1-\mu)/\mu$:

$$\mathcal{A}(\mu) = \delta^{n+1} + \frac{c_n}{r} = \left(\delta^{n+1} - c_n\right) + \frac{c_n}{\mu}, \qquad c_n \equiv \sum_{k=1}^{n}\left(\delta^k - \delta^{n+1}\right) = \frac{\delta}{1-\delta}\left[(1-\delta^n) - n(1-\delta)\delta^n\right] \geq 0,$$

where $c_n > 0$ for $n \geq 1$ by Step 3 and $c_0 = 0$. Hence, between kinks, $\mathcal{A}_{\mu\mu} = 2c_n/\mu^3 > 0$: $\mathcal{A}$ is strictly convex whenever $\mu > 1/2$ (and constant at δ for $\mu < 1/2$).

Output per order, however, is linear between kinks:

$$\mathcal{Y}(\mu) = (1-\mu) + \mu\mathcal{A}(\mu) = (1 + c_n) - s_n\mu, \qquad s_n \equiv 1 + c_n - \delta^{n+1}.$$

Since $c_{n+1} - c_n = (n+1)\delta^{n+1}(1-\delta)$, the slope changes at each kink by

$$s_{n+1} - s_n = (c_{n+1} - c_n) - (\delta^{n+2} - \delta^{n+1}) = (n+2)\delta^{n+1}(1-\delta) > 0,$$

so the slope $-s_n$ falls at each kink. Since $\mathcal{Y}$ is continuous (Step 1), piecewise linear, and has decreasing slopes, it is concave on $(0,1)$ (weakly, as it is linear between kinks). The same argument shows that $\mathcal{A}$'s slope, $-c_n/\mu^2$, falls at each kink, so the kinks of $\mathcal{A}$ are concave even though $\mathcal{A}$ is convex in between.

Step 5: Convexity in δ. For any μ, the terms $\{p, n\}$, are fixed. The second derivative of δ^a is $(a)(a-1)\delta^a$ which is strictly positive for any $a > 1$.

□

A.3 Related Results Used Elsewhere

In this section, we derived properties that are used later in the text.

Part 5. Inverse productivity. Now, we study the inverse of productivity. Let

$$q(\mu;\delta) = \mathcal{A}^{-1}(\mu;\delta).$$

The function has limits:

$$\lim_{\mu\to 0} q(\mu;\delta) = \delta^{-1}, \qquad \lim_{\mu\to 1} q(\mu;\delta) = \infty, \qquad \lim_{\delta\to 0} q(\mu;\delta) = \infty, \qquad \lim_{\delta\to 1} q(\mu;\delta) = 1.$$

Under the random assignment, the convexity follows from the concavity of $\mathcal{A}$ and the fact that it takes positive values:

$$q_{\mu\mu} = -\frac{\mathcal{A}_{\mu\mu}}{\mathcal{A}^2} + 2\frac{\mathcal{A}_\mu^2}{\mathcal{A}^3} > 0.$$

Hence, q is convex in μ.

Part 6. Elasticity of $\mathcal{A}$. A useful object in later derivations is the elasticity of $\mathcal{A}$ with respect to μ:

$$\epsilon_\mu^{\mathcal{A}} \equiv \frac{\mathcal{A}_\mu(\mu)\,\mu}{\mathcal{A}(\mu)}.$$

Using (18), under the random assignment, we obtain:

$$\mathcal{A}_\mu = -\frac{\delta(1-\delta)}{(1-\delta\mu)^2}$$

$$\frac{\mathcal{A}_\mu(\mu)\,\mu}{\mathcal{A}(\mu)} = -\frac{(1-\delta)\mu}{(1-\delta\mu)(1-\mu)} < 0.$$

Recall the relationship,

$$\frac{\partial}{\partial\mu}[\mathcal{A}(\mu)\,\mu] = \mathcal{A}'(\mu)\,\mu + \mathcal{A}(\mu) = \mathcal{A}(\mu)\,(1+\epsilon_\mu^{\mathcal{A}}).$$

We have:

$$1+\epsilon_\mu^{\mathcal{A}} = \frac{(1-\mu)(1-\delta\mu)-(1-\delta)\mu}{(1-\delta\mu)(1-\mu)} = \frac{1-\mu+\delta\mu^2-\mu}{(1-\delta\mu)(1-\mu)} = \frac{1}{1-\delta\mu} - \frac{\mu}{1-\mu}.$$

In later proofs, we require obtaining the limits of $1+\epsilon_\mu^{\mathcal{A}}$ and establishing monotonicity.

The limits of the function that governs the sign are:

$$\lim_{\mu\to 0}\left[\frac{1}{1-\delta\mu} - \frac{\mu}{1-\mu}\right] = 1 > 0, \text{ and } \lim_{\mu\to 1}\left[\frac{1}{1-\delta\mu} - \frac{\mu}{1-\mu}\right] = -\infty.$$

Since $\epsilon_\mu^{\mathcal{A}}$ is continuous in μ, the sign is ambiguous.

Next, to establish monotonicity, notice that

$$\frac{\partial}{\partial\mu}\left[\frac{1}{1-\delta\mu} - \frac{\mu}{1-\mu}\right] = \frac{\delta}{(1-\delta\mu)^2} - \frac{1}{(1-\mu)^2} < 0.$$

The numerator of the first term is lower than that of the second, whereas the opposite is true about the denominators. Hence, the function is monotone decreasing. Thus, there is a unique crossing point where the function $1+\epsilon_\mu^{\mathcal{A}}$ becomes negative.

B Extensions of the Payment-Chain Block

Economic Mechanisms for Network Formation. Random and efficient assignments analyzed above can arise from different market structures and environments. Random assignments arise naturally in decentralized markets, where, as we show in Section 4, a Walrasian mechanism equates supply and demand for inputs but does not organize orders. In such a market, agents placing chained orders trade with uncertainty about the timing of inflows and the delivery of goods. This randomness may also result from the randomness of who wants to buy from whom. In turn, an efficient assignment emerges if there exists a market entity or mechanism that can organize orders and payments, thereby minimizing payment times. Section 2.2 discusses how an entity that centralizes payments aggregates as a production function. This assignment requires substantial information and coordination, which is not necessarily the most realistic assumption.

Alternative Assignments and Empirical Patterns. Evidence from the timing-based classification of firms (Appendix C.1.3 and Figure C.4) suggests that real payment networks lie between the polar extremes of random and efficient networks: The observed distribution is bell-shaped and centered on chains of length 4–5, in contrast to the geometric distribution implied by randomness. The empirical distribution also shows a substantial share of zero-length, which indicates inefficiency.[52] These patterns motivate the alternative assignments studied in this appendix.

A first extension explores **correlated assignments** (Section B.1), where links between orders of different types do not have i.i.d. probabilities of being linked. Instead, the likelihood of a link between chained orders is $\mu + \rho$ leading to a geometric distribution of chain lengths with parameter $\mu + \rho$, and efficiency $\mathcal{A}(\mu; \delta, \rho) = \delta \frac{1-(\mu+\rho)}{1-\delta(\mu+\rho)}$.[53]

We then consider a **pooled assignment** (Section B.2), a hybrid between random and efficient allocations where chained orders can be reorganized efficiently between chains, but spot-to-spot links cannot. The formula for $\mathcal{A}$ matches the efficient assignment but replaces the ratio $\mu/(1-\mu)$ with $1/(1-\mu)$, the ratio of chained to available spot orders. This pooled assignment captures income pooling when the same agent places many chained orders. The relevance of this assignment becomes clear when we aggregate budgets in Section 2.2.

[52] Recall that the efficient size distribution is concentrated in two contiguous integers.

[53] We show this by exploiting that the payment-chain network has a Markov chain representation. When chained orders are less likely to connect with chained orders, $\rho < 0$, the economy achieves greater efficiency than under baseline i.i.d. random assignment, though falling short of full efficiency.

Another possibility is costly **reassignment** (Section B.3), another hybrid between random and efficient chains. This assignment arises when orders are initially assigned randomly, but chain orders can be reallocated at a cost Ξ, thought of as negotiation costs. A threshold position η emerges endogenously from the trade-off between reassignment costs and losses from a position. The resulting distribution exhibits efficient clustering at some position m and then a geometric distribution up to η.

Variations to the Environment. We consider three further extensions that do not modify the assignment, but the output function y_n. In one of these extensions, we allow for **default and cascading** (Section B.4), where orders at position $n > \eta$ default, keeping funds rather than executing payments. This captures situations where the losses from delays exceed the opportunity cost of not obtaining goods. Once a transaction fails, subsequent transactions also fail due to a lack of funds. In this case, $y_n = 0$ for positions above the default threshold.[54] This extension is interesting because, unlike the reassignment extension, which requires information, defaulting is a unilateral decision that only requires knowing the passage of time and can also increase efficiency.

Finally, we consider **delay heterogeneity** (Section B.5). In this extension, within any chain of length n, a number k of transactions is delayed, where k is drawn from a binomial with parameter ω and n draws, and the $n-k$ immediate transactions precede the delayed ones. The resulting TFP formula decomposes into:

$$\mathcal{A}_\omega = (1-\omega) + \delta \frac{\omega(1-\mu)}{1-\mu(1-\omega+\delta\omega)}. \tag{21}$$

When $\omega = 0$ (all instant), $\mathcal{A}_\omega = 1$; when $\omega = 1$ (all delayed), we recover the baseline formula.[55] This extension captures payment environments where instant payments coexist with traditional systems, e.g., checks. The extension is also meant to capture the possibility of distinguishing between trade credit and spot transactions along a supply chain (supply-chain finance).[56]

Finally, we generalize the formula for y_n to allow for **alternative uses** of resources while products are not delivered. In the microfounded setting for y_n, we assume that

[54] If the chain-length distribution remains geometric, the TFP is the baseline TFP scaled by $(1-\delta^\eta p^\eta)$, where $p = \mu + \rho$ is the probability that a chain continues ($p = \mu$ under the random assignment)

[55] The first term $(1-\omega)$ is the contribution from instant payments, while the second term captures the reduced productivity from delayed transactions.

[56] We treat spot orders as including trade credit. However, there is a nuanced distinction when the supplier extending trade credit is financially constrained. In such cases, trade credit ceases to be equivalent to a spot transaction and becomes a hybrid. If the downstream firm lacks liquidity, it must wait for the upstream firm to sell before repaying its obligations. Trade credit thus accelerates production at one node but delays another. The heterogeneous chain captures this possibility. Bocola (2022) studies cases where buyers and sellers are constrained, but does not feature the delays emphasized in our extension.

labor remains idle until inputs are delivered. Suppose instead that labor has a less efficient alternative use, producing $(1-\alpha)$ units of final goods per unit of time *before* inputs arrive and one unit of final goods per unit of time *after* inputs arrive at τ_n. Then, y_n is a weighted average of the outputs per unit of time, weighted by the duration when inputs are unavailable versus available:

$$y_n = (1-\alpha)(1-\delta^n) + \delta^n = 1 - \alpha + \alpha\delta^n.$$

Using the same aggregation properties and expectations' linearity, aggregate output is:

$$\mathcal{Y} = (1-\mu)\cdot 1 + \mu \cdot \mathbb{E}[1-\alpha+\alpha\delta^n | n \geq 1] = (1-\alpha\mu) + \mu\alpha \cdot \mathcal{A}(\mu,\delta).$$

The parameter $\alpha \in [0,1]$ captures the productivity loss from not having the specialized input: $\alpha = 1$ recovers our baseline (complete idleness), while $\alpha = 0$ implies no productivity loss (perfect substitutes).

B.1 Correlated Assignments

In the main text, we considered payment chains where spot and chained orders have equal probabilities of being linked to orders of the same type. However, the timing-based classification of firms in Table C.6 reveals that payment types exhibit significant correlation: chained orders predominantly connect to other chained orders, while spot orders tend to cluster together. This correlation in order types affects the distribution of chain lengths and ultimately productivity. To examine this, we allow the probability that chained orders link to other chained orders to differ from the probability that spot orders connect to chained orders. We formalize this using a Markov chain representation.

Markov Chain Representation. Consider a payment-chain network where order types (spot or chained) exhibit correlation along production chains. To ease the notation, let $\mu_x = \mu$ denote the fraction of chained orders and $\mu_s = 1 - \mu_x$ the fraction of spot orders. Starting from a spot order, the entire payment network is characterized by a sequence $S_t = \{s, x, s, x, s, \ldots\}$ where $\{s, x\}$ in position i represents whether the i-th order is spot or chained, respectively. If links across chained and spot orders only depend on the current order type, the sequence can be generated by a Markov chain:

$$Q = \begin{bmatrix} p_{ss} & 1-p_{xx} \\ 1-p_{ss} & p_{xx} \end{bmatrix} \tag{22}$$

where p_{ss} is the probability that a spot order is linked to a spot, and p_{xx} the corresponding probability that a chained order is linked to a chained order.

For a large N, the stationary distribution must satisfy:

$$\boldsymbol{\mu} = Q\boldsymbol{\mu}, \quad \text{where } \boldsymbol{\mu} = \begin{bmatrix} \mu_s \\ \mu_x \end{bmatrix}. \tag{23}$$

That is, the transition matrix must guarantee that population shares aggregate.

To see how the construction works, we verify the formulas for the i.i.d. case in which transition probabilities equal unconditional probabilities:

$$Q^{\text{iid}} = \begin{bmatrix} \mu_s & \mu_s \\ \mu_x & \mu_x \end{bmatrix}. \tag{24}$$

This trivially satisfies $\boldsymbol{\mu} = Q^{\text{iid}}\boldsymbol{\mu}$ since each column sums to one.

To introduce correlation, we parameterize deviations from independence as:

$$p_{ss} = \mu_s\kappa_s, \quad p_{xx} = \mu_x\kappa_x \tag{25}$$

where $\kappa_s, \kappa_x > 0$ are correlation factors. When the corresponding κ exceeds one, orders have a greater chance of matching with orders of the same type than under an i.i.d. assignment; when κ is below one, the random assignment overrepresents mixed links.

For positive values of κ, the transition matrix becomes:

$$Q = Q^{\text{iid}} + \tilde{Q} \tag{26}$$

where the deviation matrix is:

$$\tilde{Q} = \begin{bmatrix} \mu_s(\kappa_s - 1) & -\mu_x(\kappa_x - 1) \\ -\mu_s(\kappa_s - 1) & \mu_x(\kappa_x - 1) \end{bmatrix}. \tag{27}$$

There are some parameter constraints: We require a stationarity condition on $\tilde{Q}\boldsymbol{\mu} = 0$ that guarantees the preservation of population shares:

$$\mu_s^2(\kappa_s - 1) = \mu_x^2(\kappa_x - 1). \tag{28}$$

Defining $\rho_s \equiv \mu_s(\kappa_s - 1)$ and using the constraint, we obtain:

$$\mu_x(\kappa_x - 1) = \frac{\mu_s}{\mu_x}\rho_s. \tag{29}$$

Setting $\rho \equiv \frac{\mu_s}{\mu_x}\rho_s$, the transition matrix simplifies to:

$$Q(\rho) = \begin{bmatrix} \mu_s + \rho_s & 1 - (\mu_x + \rho) \\ 1 - (\mu_s + \rho_s) & \mu_x + \rho \end{bmatrix}. \tag{30}$$

To ensure valid probabilities ($0 \leq p_{ij} \leq 1$), we require: if $\mu_s > 1/2$, then $-\mu_s \leq \rho_s \leq 1 - \mu_s$; if $\mu_x > 1/2$, then $-\mu_x \leq \frac{\mu_s}{\mu_x}\rho_s \leq 1 - \mu_x$. Combining these constraints:

$$-\min\{\mu_x, \mu_s\} \leq \rho \leq \min\{1 - \mu_s, 1 - \mu_x\}. \tag{31}$$

Chain Length Distribution. Conditional on featuring some chained orders, the length of a chain (consecutive chained orders) follows a geometric distribution. A chain of length n occurs with probability:

$$\Pr(n|n \geq 1) = P(\text{chain ends}) \cdot P(\text{continue chain})^{n-1}. \tag{32}$$

Given the parameterization:

$$P(\text{continue chain}) = p_{xx} = \mu_x + \rho \tag{33}$$

$$P(\text{chain ends}) = 1 - p_{xx} = 1 - \mu_x - \rho \tag{34}$$

Therefore:

$$\Pr(n|n \geq 1) = (1 - \mu_x - \rho)(\mu_x + \rho)^{n-1}. \tag{35}$$

This is a geometric distribution with success probability $p = \mu_x + \rho$.

The expected chain length conditional on some chained orders is:

$$\mathbb{E}[n|n \geq 1] = \frac{1}{1-p} = \frac{1}{1 - \mu_x - \rho}. \tag{36}$$

Thus, the unconditional distribution is:

$$\mathbb{E}[n] = \mathbb{E}[n|n \geq 1]\Pr[n \geq 1] = \frac{\mu_x - \rho_s}{1 - \mu_x - \rho} = \frac{\mu_x - \frac{\mu_x}{\mu_s}\rho}{1 - \mu_x - \rho},$$

where we used that $\rho_s \equiv \frac{\mu_x}{\mu_s}\rho$. Notice that when $\rho > 0$, we have positive correlation—chained orders cluster together, increasing average chain length; when $\rho < 0$, the correlation decreases average chain length. When $\rho = 0$ (independence), we recover $\mathbb{E}[L] = \frac{\mu}{1-\mu}$.

Output and TFP. Using the same steps as we do for the i.i.d. case, using the chain-length distribution method, we obtain:

$$\mathcal{A}(\mu; \delta, \rho) = \delta \frac{1 - (\mu_x + \rho)}{1 - \delta(\mu_x + \rho)}. \tag{37}$$

Output is computed by weighting this distribution:

$$\mathcal{Y}(\mu; \delta, \rho) = (1 - \mu) + \mu\mathcal{A} = 1 - \mu \frac{1 - \delta}{1 - \delta(\mu + \rho)}. \tag{38}$$

Again, this verifies the consistency of the formulas when $\rho = 0$.

B.2 Pooled Chained Orders

Next, we consider assignments where chained orders can be efficiently reassigned to chains of approximately equal length, while the mass of spot-to-spot orders remains unchanged. Under this mechanism, the income from every chained order is immediately used to finance another chained order as long as unfunded chained orders exist. This captures scenarios where agents placing multiple chained orders pool their funds, ensuring no cash remains idle while chained orders await funding.

Setup with Correlation. As with correlated assignments, the probability that a spot order is linked to a chained order is $(1 - \mu_s - \rho_s)$, where ρ_s is the spot-to-spot correlation parameter from our analysis above. While spot-to-spot matches follow a random process, any spot order that connects to a chained order enters a pool where chains are optimally formed to maximize total production. Thus, the mass of spot orders that can potentially match with chained orders is constant:

$$\bar{N}_s = (1 - \mu)(1 - \mu_s - \rho_s)N = (1 - \mu)(\mu - \rho_s)N. \tag{39}$$

The ratio of chained orders to available spot orders is:

$$\xi(\rho_s) = \frac{N_x}{\bar{N}_s} = \frac{\mu N}{(1-\mu)(\mu-\rho_s)N} = \frac{\mu}{(1-\mu)(\mu-\rho_s)} \tag{40}$$

When $\rho_s > 0$ (positive correlation): Fewer spot orders connect to chained orders, increasing ξ; when $\rho_s < 0$ (negative correlation): More spot orders connect to chained orders, decreasing ξ toward the efficient ratio. When $\rho_s = 0$ (independence), we obtain $\xi(0) = \frac{1}{1-\mu}$, thus the ratio of chained to spot orders exceeds $\frac{\mu}{1-\mu}$. This reveals that the assignment is less efficient than the fully efficient case, since some spot orders are inefficiently matched with other spot orders. When $\rho_s = -(1-\mu)$, we have $\xi(-(1-\mu)) = \frac{\mu}{1-\mu}$, recovering the fully efficient assignment case.

Optimal Chain Length Distribution. Following the arguments of the efficient case, the efficient assignment creates chains of two adjacent lengths for chains of length greater than or equal to one:

$$n(\rho_s) = \lfloor \xi(\rho_s) \rfloor \tag{41}$$

$$n(\rho_s) + 1 = \lfloor \xi(\rho_s) \rfloor + 1 \tag{42}$$

The fraction of chains with the shorter length is:

$$p(\rho_s) = \lfloor \xi(\rho_s) \rfloor + 1 - \xi(\rho_s) \tag{43}$$

The resulting distribution is:

$$g(0) = 1 - (\mu - \rho_s) \quad \text{(spot-only transactions)} \tag{44}$$

$$g(n(\rho_s)) = p(\rho_s)(\mu - \rho_s) \quad \text{(shorter chains)} \tag{45}$$

$$g(n(\rho_s)+1) = (1 - p(\rho_s))(\mu - \rho_s) \quad \text{(longer chains)} \tag{46}$$

Clearly, the pooling mechanism partially corrects the inefficiency from random matching, but maintains the inefficiencies provoked by the assignment of spot-to-spot links.

Production and TFP. TFP is in this case given by:

$$\mathcal{A}(\mu, \rho_s) = \frac{p(\rho_s) \cdot Y(n(\rho_s)) + (1 - p(\rho_s)) \cdot Y(n(\rho_s)+1)}{p(\rho_s) \cdot n(\rho_s) + (1 - p(\rho_s)) \cdot (n(\rho_s)+1)} \tag{47}$$

where $Y(n) = \delta\frac{1-\delta^n}{1-\delta}$. Expected production per order follows the baseline formula, $\mathcal{Y}(\mu,\rho_s) = (1-\mu) + \mu\mathcal{A}(\mu,\rho_s)$.

B.3 Endogenous Reallocation

The pooled assignment above assumes costless reallocation of funds across chained orders. However, it is natural to expect that reallocation only occurs for orders in positions exceeding a certain threshold. We now examine a model where chains beyond a threshold length η must reallocate their excess orders to shorter chains, while chains below this threshold operate normally.

Setup. Consider a payment network where chains longer than the threshold η must reallocate their excess orders. We define the **orders needing reallocation** (N^r) as the total number of chained orders in positions beyond η:

$$N^r \equiv \sum_{n=\eta+1}^{\infty} (n-\eta)\cdot g(n) = \frac{\mu^{\eta+1}}{1-\mu} \tag{48}$$

This counts all orders in chains of length $n > \eta$ that exceed the threshold.

The **available slots** (N^a) in other chains is the total unused capacity in chains of length $n \leq \eta$:

$$N^a \equiv \sum_{n=0}^{\eta}(\eta-n)\cdot g(n) = \eta - \left(\frac{\eta\mu^{\eta+1}}{1-\mu}\right)(2-\mu) - \frac{1}{1-\mu} + \frac{(1+\eta)\mu^{\eta}}{1-\mu} \tag{49}$$

Each chain of length $n \leq \eta$ can potentially accept $(\eta - n)$ additional orders before reaching the threshold.

The ratio of reallocating orders to available slots determines the efficient assignment structure:

$$\xi(\eta) = \frac{N^r}{N^a} = \frac{\mu^{\eta}}{(\eta-1)(1-\mu) + \mu^{\eta}[1+\eta(1-\mu)^2]} \tag{50}$$

This ratio determines the optimal chain lengths for efficient assignment:

$$n(\eta) = \lfloor \xi(\eta) \rfloor \tag{51}$$

$$n(\eta) + 1 = \lceil \xi(\eta) \rceil \tag{52}$$

with fraction $p(\eta) = \lceil \xi(\eta) \rceil - \xi(\eta)$ assigned to the shorter length.

In this setting, orders are reallocated sequentially by approaching the customers

and suppliers in the last position of an existing chain. They offer to buy from the supplier (corresponding to a spot order) and sell to the last customer order (corresponding to a chain order). When abandoning the chain, the final customer buys from the preceding customer. The reallocation has a cost of Ξ.

Chain-Length Distribution. Because the assignment is sequential, reassigned orders will be placed at the shortest available slots, including those from spot orders. As a result, the distribution will be efficient among the reassigned orders.

Let $g(k) = (1-\mu)\mu^{k-1}$ for $k \geq 1$ denote the baseline geometric distribution for chain lengths, and $G(k) = \sum_{j=1}^{k} g(j) = 1 - \mu^k$ its cumulative distribution function. Note that in the baseline model, spot orders have mass $(1-\mu)$.

After the re-assignment, the resulting hybrid distribution for chain lengths is:

$$\hat{g}(k;\eta,\mu) = \begin{cases} \alpha \cdot p(\eta) & k = n(\eta) \\ \alpha \cdot (1 - p(\eta)) & k = n(\eta)+1 \\ g(k) & n(\eta)+2 \leq k \leq \eta - 1 \\ g(\eta) + (1 - G(\eta) & k = \eta \text{ (truncated chains)} \\ 0 & \text{otherwise} \end{cases} \tag{53}$$

where $\alpha = G(n(\eta)+1)$ is the CDF of chains that would have been length $\leq n(\eta)+1$ in the baseline. This distribution has three key features: (i) all spot orders have been converted into chains through the efficient assignment, (ii) the original distribution between lengths $n(\eta)+1$ and $\eta - 1$, is preserved reflecting that orders in those positions are not re-assigned, (iii) the mass at η equals μ^η, capturing that all chains with length above η become chains of length η.

Endogenous Choice. Given the cost Ξ and the sequential nature of the re-assignment, the value of η can be found by solving for the value for which:

$$\Xi < \delta^{n(\eta)+1} - \delta^{\eta+1}, \qquad \Xi \geq \delta^{n(\eta)+1} - \delta^{\eta}.$$

The left inequality ensures reallocation from position $\eta+1$ to the last position with positive mass $n(\eta)+1$ is profitable, while the right inequality ensures reallocation from position η is not.

B.4 Default and Cascading

The reassignment model above assumes agents can identify and relocate to shorter chains. However, this requires substantial information about the network. We now consider a simpler response: when chains exceed a threshold η, orders beyond this point are withdrawn rather than reassigned. This captures situations where extended payment chains become unsustainable. Orders in positions $n > \eta$ simply default, keeping their funds. When they do so, all of the subsequent transactions experience a cascading sequence of breakdowns.

Chain Length Distribution. As before, chain length follows a geometric distribution:

$$g(n) = (1-p)p^{n-1} \tag{54}$$

where p generalizes to include correlation effects: $p = \mu + \rho$ in the correlated case.

We capture default by modifying the output function to:

$$y_n = \begin{cases} \delta^n & \text{if } n \leq \eta \\ 0 & \text{if } n > \eta. \end{cases} \tag{55}$$

The idea is that chained orders withdraw their orders whenever they fall in position $n > \eta$. In that case, production among chained orders in position below or equal to η remains the same, but orders in higher position break down and obtain zero output.

In that case, the average output of chains of length n is given by:

$$\bar{y}(n) = \begin{cases} \frac{\delta}{n}\frac{1-\delta^n}{1-\delta} & \text{if } n \leq \eta \\ \frac{\delta}{n}\left(\frac{1-\delta^\eta}{1-\delta}\right) & \text{if } n > \eta. \end{cases} \tag{56}$$

Average production, thus, can be written as:

$$\bar{y}(n) = \frac{y(n)}{n} = \frac{\delta}{n}\frac{1-\delta^n}{1-\delta} - \bar{\gamma}(n) \tag{57}$$

where the default adjustment is:

$$\bar{\gamma}(n) = \begin{cases} 0 & \text{if } n \leq \eta \\ \frac{1}{n}\frac{\delta^\eta}{1-\delta}(1-\delta^{n-\eta}) & \text{if } n > \eta. \end{cases} \tag{58}$$

Using the methods presented for the baseline case, the formula for TFP becomes:

$$\mathcal{A}_\eta = \frac{\sum_{n=1}^{\infty} g(n)\bar{y}(n)}{\sum_{n=1}^{\infty} g(n)n}. \tag{59}$$

Combining all terms we obtain that:

$$\mathcal{A}_\eta = \frac{\delta(1-p)}{1-p\delta} - \frac{\delta(1-p)}{1-p\delta}\delta^\eta p^\eta = \mathcal{A}(1-\delta^\eta p^\eta) \tag{60}$$

where $\mathcal{A} = \frac{\delta(1-p)}{1-p\delta}$ is the TFP without default. The truncation factor $(1-\delta^\eta p^\eta)$ captures the productivity loss from default. This loss increases the lower threshold η (more chains are truncated). When $\eta \to \infty$, we recover $\mathcal{A}_\eta \to \mathcal{A}$, the case without default. When $\eta = 0$, note that all chained orders default, leading to zero output among chained orders. Output follows the formula for the correlated case, replacing $\mathcal{A}$ by $\mathcal{A}_\eta$.

B.5 Partial Delays

We now consider a hybrid payment system in which each transaction in a chain has a probability ω of occurring with a delay. This captures heterogeneity in payment processing—some transactions use sophisticated instant payment systems while others require traditional settlement with delays.

Modified Production Structure. For a chain of length n where k transactions happen, production depends critically on the *location* of delays. We assume the first $n-k$ links are paid immediately and the last k are delayed, yielding:

$$y(k,n) = (n-k) + \delta\frac{1-\delta^k}{1-\delta} \tag{61}$$

where $(n-k)$ represents immediate payments and the second term captures delayed payments.

Distribution of Delays and TFP. The number of delayed payments in a chain of length n follows a binomial distribution:

$$M(k,n) = \binom{n}{k}\omega^k(1-\omega)^{n-k}. \tag{62}$$

The expected production for a chain of length n is:

$$\mathbb{E}[y(k,n)|n] = \sum_{k=0}^{n} y(k,n) \times M(k,n). \tag{63}$$

To evaluate this expectation, we use three key properties of the binomial distribution:

$$\sum_{k=0}^{n} M(k,n) = 1 \tag{64}$$

$$\sum_{k=0}^{n} k \cdot M(k,n) = qn \tag{65}$$

$$\sum_{k=0}^{n} \delta^k M(k,n) = (1-\omega+\delta\omega)^n. \tag{66}$$

The third property follows from the binomial theorem.[57] Substituting these properties:

$$\mathbb{E}[y(k,n)|n] = \sum_{k=0}^{n}(n-k)M(k,n) + \frac{\delta}{1-\delta}\left[1 - \sum_{k=0}^{n}\delta^k M(k,n)\right] \tag{68}$$

$$= (1-\omega)n + \frac{\delta}{1-\delta}\left[1-(1-\omega+\delta\omega)^n\right]. \tag{69}$$

Average production across all chain lengths, weighted by the geometric distribution $g(n) = (1-p)p^{n-1}$ is given by:

$$\mathcal{A}_\omega = \frac{\sum_{n=1}^{\infty} g(n) \cdot \mathbb{E}[y(k,n)|n]}{\sum_{n=1}^{\infty} g(n) \cdot n} \tag{70}$$

The denominator equals $\frac{1}{1-p}$. For the numerator:

$$(1-p)\sum_{n=1}^{\infty} p^{n-1}\mathbb{E}[y(k,n)|n] \tag{71}$$

$$= \frac{1-\omega}{1-p} + \frac{\delta}{1-\delta}\left[1-(1-p)\frac{1-\omega+\delta\omega}{1-p(1-\omega+\delta\omega)}\right] \tag{72}$$

[57] The assumption that immediate payments precede delayed ones is crucial. If delays were randomly distributed throughout the chain, the production function would be more complex, as non-delayed segments would be interrupted by delayed ones. Since a payment that follows a delayed payment inherits its delay, ordering immediate payments first ensures that positions 1 through $n-k$ are realized without delay. For the third property, recall that

$$\sum_{k=0}^{n}\binom{n}{k}(\delta\omega)^k(1-\omega)^{n-k} = (1-\omega+\delta\omega)^n. \tag{67}$$

Simplifying yields:

$$\mathcal{A}_\omega = (1-\omega) + \delta \frac{\omega(1-p)}{1-p(1-\omega+\delta\omega)} \tag{73}$$

This generalizes our baseline formula: When $\omega = 1$ (full delays): $\mathcal{A}_\omega = \frac{\delta(1-p)}{1-p\delta}$. When $\omega = 0$ (no delays): $\mathcal{A}_\omega = 1$ When $p = \mu$ and $\omega = 1$, it recovers the baseline uncorrelated case.

C Appendix to the Empirical Analysis (Section 3)

C.1 Validation: Incoming liquidity predicts outgoing payments

Chained orders arise when firms use incoming payments to finance outgoing payments. Thus, a payment received today should predict a payment sent shortly thereafter, and the relationship should be strongest for firms with tighter liquidity. This relation is difficult to establish without timestamps and payer-payee identifiers for the universe of payments, as accounting data generally record only stocks or low-frequency flows.

To gauge how often outgoing payments are tied to the arrival of incoming liquidity, we estimate

$$sent_{jt} = \beta\, received_{j,t-1} + \lambda_t + \lambda_{j\times d} + \varepsilon_{jt}, \tag{74}$$

where $sent_{jt} = 1$ if firm j made a payment on business day t, and $received_{j,t-1} = 1$ if firm j received a payment in the preceding 24 hours. The fixed effects make the comparison demanding. Calendar-time fixed effects, λ_t, absorb economy-wide payment cycles and shocks common to all firms on a given day, such as month-end settlement, tax deadlines, holidays, banking disruptions, or aggregate demand conditions. Firm-by-day-of-week fixed effects, $\lambda_{j\times d}$, absorb each firm's normal payment schedule: for example, that a particular firm usually pays suppliers on Mondays, or that another firm rarely makes payments on Fridays. Thus, β is identified from deviations from a firm's own usual payment behavior on that weekday, comparing days that follow an incoming payment to otherwise similar days that do not. A positive coefficient therefore means that, after receiving money, a firm is unusually likely to send money out, over and above both aggregate payment cycles and the firm's own regular payment routine.

Table C.1 reports the results. Column (1) shows that, on average, firms are significantly more likely to send a payment within 24 hours of receiving a payment. Among realized receipt–payment pairs, the association is concentrated when the amount sent is smaller than the amount received, as shown in Table C.2. Columns (2)–(4) reveal heterogeneity: the payment–chain effect is strongest among small firms (5 workers or less) and attenuated for firms with higher asset-to-sales ratios or higher liabilities-to-sales ratios.

Columns (1)–(4) focus on a one-day window to approximate chained-payment responses, as a narrow time frame strengthens the inference of a direct link between inflows and outflows. To trace the temporal profile of this effect, we implement a local projection (Jordà, 2005; Jordà and Taylor, 2025). Figure C.1 shows an immediate response within one business day and a second local peak around five business days

Table C.1: Triggered Payments Across Firms

Dependent variable: Probability that firm j makes a payment at time t

	(1)	(2)	(3)	(4)	(5)	(6)
$received_{j,t-1}$	0.030	0.010	0.039	0.038		0.030
	(0.0005)***	(0.0008)***	(0.0008)***	(0.0008)***		(0.0005)***
$Small_j \times received_{j,t-1}$		0.017				
		(0.001)***				
$\frac{Assets_j}{Sales_j} \times received_{j,t-1}$			-0.019			
			(0.001)***			
$\frac{Liabilities_j}{Sales_j} \times received_{j,t-1}$				-0.018		
				(0.001)***		
$received_{j,t-5}$					0.011	0.010
					(0.0004)***	(0.0004)***
Adjusted R^2	0.483	0.446	0.483	0.483	0.483	0.483
Observations	10,001,248	5,290,408	8,399,599	8,399,599	10,001,248	10,001,248
Clusters	222,306	74,566	174,896	174,896	222,306	222,306
$t, j \times$ d FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The independent variable in column (1) is an indicator that equals one if firm j received a payment from another firm in the last 24 hours. This variable is interacted with firm j's size, assets-to-sales, and liabilities-to-sales in columns (2), (3), and (4), respectively. Column (5)'s independent variable is an indicator that equals one if firm j received a payment exactly five business days before t. Column (6) includes both the next-day and five-business-day indicators. Robust standard errors, adjusted for clustering by firm, are in parentheses. We include time and firm×weekday fixed-effects. Data are daily and span 2024.

after receipt of funds. This two-wave structure reflects both the speed of Costa Rica's modern payment infrastructure and residual frictions. In 2024, when Sinpe processes most transfers electronically, next-day responses are technologically feasible.[58] However, even with a rapid payments network in place, a subset of firms continues to rely on paper checks, which take several business days to clear. This check-clearance lag produces the pronounced second spike at business day five. Consistent with this mechanism, firms exhibiting the five-day payment lag are roughly 60% more likely to pay by check than the rest. Moreover, when we exclude these check—using firms, the day-five peak of Figure C.1 disappears, confirming that check—clearing times underlie the longer delays. Coefficients in columns (1), (5), and (6) are remarkably stable, suggesting that payments unfold in two orthogonal waves—an immediate response and a delayed one around business day five.

[58] By contrast, in the 1990s, most payments settled by check, and checks used to require up to 15 days to clear—so that the fastest possible response was about two weeks (Cerdas and Melegatti, 2014).

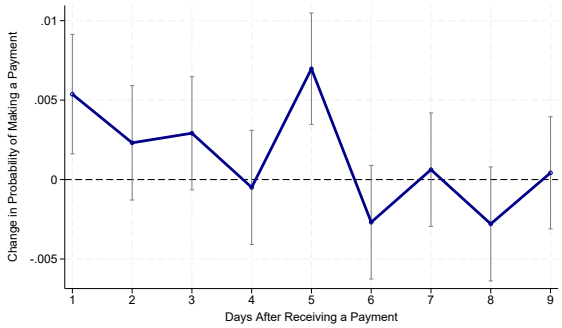


Figure C.1: Triggered Payments: Local Projection

Notes: The figure shows a local projection (Jordà, 2005; Jordà and Taylor, 2025), where the dependent and independent variables have 13 lags, as determined by a VAR-based optimal lag selection criteria. Observations are weighted by sales. Data are monthly and span 2024.

Table C.2 shows that, reassuringly, the pattern is concentrated in cases where the amount sent is smaller than the amount just received, as one would expect if incoming liquidity helps finance outgoing payments. Appendix C.1.1 also shows similar results on the intensive margin and, importantly, that placebo exercises using payments received *right after* outgoing payments do *not* predict the positive increase in outgoing payments observed after past receipts; the placebo coefficients are small and nonpositive.

Taken together, these results are consistent with the payment-chain margin emphasized by the model: recent receipts systematically predict subsequent payments, especially in cases where the incoming payment can plausibly finance the outgoing one.

C.1.1 Intensive Margin Regressions

We also study the intensive margin relationship between payments made and received. We follow equation (74), but now considering the (inverse hyperbolic sine of the) amounts sent and received by firm j. Results are presented in Table C.3 consistent with those of the extensive margins; firms become more likely to make payments once they receive money transfers. The asinh-on-asinh coefficient in column (1) is 0.036, indicating a positive intensive-margin relationship between amounts received and amounts sent the following day.

Table C.2: Heterogeneous Effect by Size of Amounts Sent/Received

Dependent variable: Prob. firm j makes a payment at time t

	(1)
$received_{j,t-1} \times smaller\ sent_{jt}$	0.526
	(0.002)***
$received_{j,t-1}$	-0.109
	(0.001)***
Adjusted R^2	0.535
Observations	10,001,248
clusters	222,306
$t, j \times$ d FE	Yes

Notes: The independent variable $received_{j,t}$ is an indicator equal to one if firm j received a transfer from another firm within the last 24 hours. This variable is interacted with $smaller\ sent_{jt}$, which equals one if the amount sent by firm j at time t is smaller than the amount received by firm j at time $t-1$, and is zero otherwise. Robust standard errors, clustered by firm, are in parentheses. We include time and firm×weekday fixed-effects. Data are daily and span 2024.

Table C.3: Triggered Payments Across Firms—Intensive Margin

Dependent variable: Payments made by firm j at time t (asinh)

	(1)	(2)	(3)	(4)	(5)	(6)
$\ln amount\ received_{j,t-1}$	0.036	0.017	0.045	0.044		0.036
(*asinh*)	(0.0005)***	(0.001)***	(0.001)***	(0.001)***		(0.0005)***
$Small_j \times \ln amount\ received_{j,t-1}$		0.016				
		(0.001)***				
$\frac{Assets_j}{Sales_j} \times \ln amount\ received_{j,t-1}$			-0.019			
			(0.001)***			
$\frac{Liabilities_j}{Sales_j} \times \ln amount\ received_{j,t-1}$				-0.018		
				(0.001)***		
$\ln amount\ received_{j,t-5}$					0.012	0.012
(*asinh*)					(0.0004)***	(0.0004)***
Adjusted R^2	0.496	0.464	0.497	0.497	0.495	0.496
Observations	10,001,248	5,290,408	8,399,599	8,399,599	10,001,248	10,001,248
Clusters	222,306	74,566	174,896	174,896	222,306	222,306
$t, j \times$ d FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The independent variable in column (1) is the total amount (in colones) received by firm j from another firm within the last 24 hours, transformed using the inverse hyperbolic sine function. This variable is interacted with firm j's size, assets-to-sales and liabilities-to-sales in columns (2), (3), and (4) respectively. Columns (5) and (6) add an independent variable equal to the total amount (in colones) received by firm j from another firm 5 days before time t, transformed using the inverse hyperbolic sine function. Robust standard errors, clustered by firm, are in parentheses. We include time and firm×weekday fixed-effects. Data are daily and span 2024.

C.1.2 Placebo Exercises

Table C.4: Placebos—Receiving Payments Right After Making a Transfer

Dependent variable: Payment made by firm j at time t

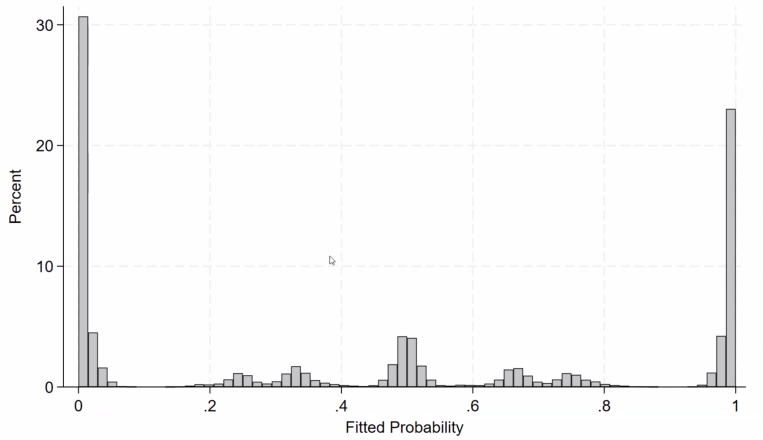
	A. Extensive margin				B. Intensive margin (*asinh*)		
	$h=1$	$h=2$	$h=3$		$h=1$	$h=2$	$h=3$
	(1)	(2)	(3)		(4)	(5)	(6)
$received_{j,t+h}$	-0.001	-0.003	-0.001	$\ln amount$	-0.001	-0.002	0.000
	(0.0004)***	(0.0004)***	(0.0004)	$received_{j,t+h}$	(0.0004)	(0.0004)***	(0.0004)
Adjusted R^2	0.483	0.483	0.483		0.495	0.495	0.495
Observations	10M	10M	10M		10M	10M	10M
Clusters	222,306	222,306	222,306		222,306	222,306	222,306
$t, j \times$ d FE	Yes	Yes	Yes		Yes	Yes	Yes

Notes: In panel A, the independent variables are indicators equal to one if firm j received a payment in $t+h$. In panel B, the independent variables are the total amounts (in colones) received by firm j from another firm in $t+h$, transformed using the inverse hyperbolic sine function. Robust standard errors, clustered by firm, are in parentheses. We include time and firm×weekday fixed-effects. Data are daily and span 2024.

C.1.3 Timing-Based Classification

Leveraging the comprehensive ACH and wire transfer data, our goal is to cluster firms into two groups: (i) firms that tend to have chained payments and (ii) firms that tend to rarely have chained payments, and instead receive spot payments. To do so, we first consider the distribution of monthly occurrences by firm when receiving a payment led to making a payment within a day. As shown in Figure C.2, there are large masses of firms concentrated in zero—which aligns with a payment behavior that is orthogonal to receiving money transfers, i.e., firms that make spot payments—and in one—which would imply a perfect correlation between receiving a transfer and making a payment the day after, i.e., firms that make chained payments. Given this largely bimodal distribution, we then use k-means clustering with $k=2$ to split the sample into two groups; firms that tend to make chained and spot payments.

Figure C.2: Histogram of Fitted Probabilities of Triggered Payments by Firm per Month



Notes: The figure shows the distribution of monthly average triggered payments by firm based on an extensive margin estimation. Firms with a value of zero have a payment behavior that is orthogonal to receiving money transfers; firms with a value of one have a perfect correlation between receiving a transfer and making a payment the day after. Data are daily and span 2024.

Correspondence with invoice and delivery timing. We next ask whether the firm-level classification based on bank-payment timing predicts the transaction-level combination of deferred payment and delayed delivery emphasized in Fact 2.

Table C.5: Correspondence Between Timing-Based and Transaction-Level Classifications

Dependent variable: deferred transaction delivered at or after the due date		
	Trade credit excluded	Trade credit coded as spot
	(1)	(2)
Timing-based chained buyer	0.071^{***}	0.058^{***}
	(0.003)	(0.002)
Month fixed effects	Yes	Yes
Dependent-variable mean	0.099	0.088
Coefficient relative to mean	71%	66%
Adjusted R^2	0.015	0.013
Observations	36,825,956	41,555,522
Buyer clusters	174,887	174,929

Notes: The dependent variable equals one when payment is deferred and the associated delivery-service payment is first observed at or after the invoice's payment deadline. The unit of observation is a transaction. In column (1), the omitted outcome consists of immediate-payment transactions; deferred transactions delivered before the due date are excluded. Column (2) instead codes both immediate-payment transactions and deferred transactions delivered before the due date as zero. The independent variable identifies buyers classified as chained from the timing of incoming and outgoing bank payments. All regressions include month fixed effects. Robust standard errors clustered at the buyer level are reported in parentheses. The data span 2024. *** denotes significance at the 1 percent level.

C.1.4 Payment Delays and Chained Behavior

In the model, longer settlement delays reduce the effective quantity obtained through chained orders and raise their effective price, creating an incentive to reduce their use. Figure C.3 shows a pattern consistent with this mechanism under two independently constructed classifications. Panel (a) uses the timing-based firm classification, while panel (b) classifies a transaction as chained when payment is deferred and the associated delivery-service payment is first observed at or after the full-payment due date. Both panels show a negative relationship between average payment delays and the incidence of chained behavior.

A. Timing-based firm classification B. Deferred payment + delivery classification

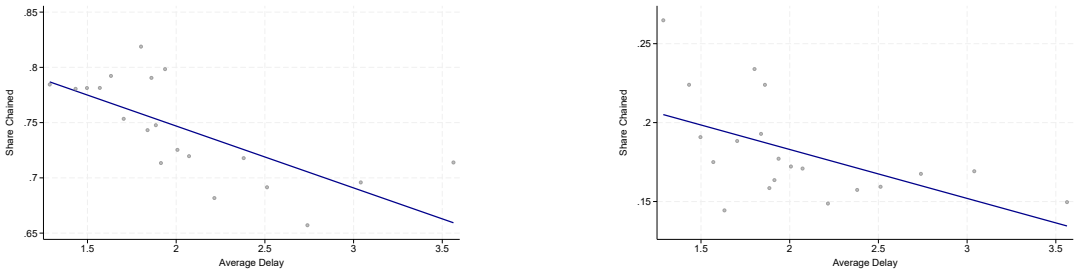
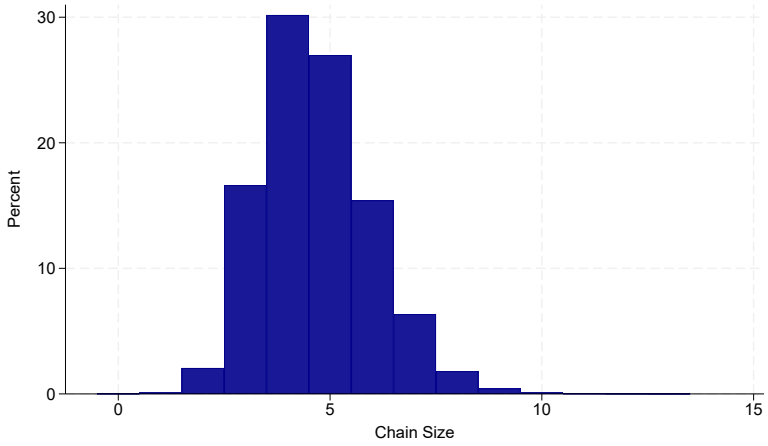


Figure C.3: Payment Delays and Chained Behavior

Notes: Panel A plots, by sector, the average share of firms classified as chained using the timing of incoming and outgoing bank payments against the sector's average payment delay. Panel B plots, by sector, the average share of transactions classified as chained using deferred payment together with delivery-service payment first observed at or after the invoice's full-payment due date against the sector's average payment delay. Sectors correspond to ISIC sections. Observations are weighted by sales. Data are monthly and span 2024.

C.1.5 Timing-Based Classification and Chains

Figure C.4: Distribution of Chain Sizes



Notes: The figure shows the distribution of chain sizes based on a 5% random sample of spot firms in the economy used as initial nodes. Data are collapsed for December 2024 to construct the chains.

Table C.6: Classification of Transactions by Month

Month	S→S	X→X	Mixed	Total	Month	S→S	X→X	Mixed	Total
1	1%	72%	28%	277,758	7	1%	72%	28%	291,866
2	1%	71%	28%	278,278	8	1%	72%	27%	276,308
3	1%	72%	28%	262,887	9	1%	72%	28%	280,320
4	1%	72%	28%	291,672	10	1%	71%	28%	290,711
5	1%	72%	28%	293,231	11	1%	71%	28%	282,262
6	1%	72%	28%	267,484	12	1%	72%	27%	287,585

Notes: The table shows the stability of the classification of transactions across months, depending on the buyer's and seller's type. Link types are divided into spot-to-spot (S→S), chained-to-chained (X→X), or mixed (S→X or X→S).

C.2 Delivery in Retail

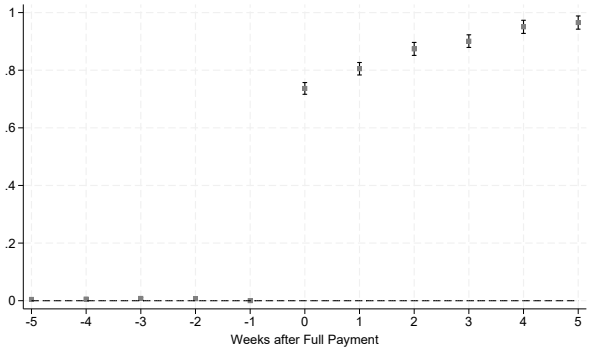


Figure C.5: Retail Sales Around Payment Deadlines

Notes: The figure includes deferred-payment purchases for which the retailer begins selling the product after its payment deadline. Time zero is the retailer's deadline for paying the supplier, and the coefficients describe the retailer's sales of the product. Error bars denote 95% confidence intervals.

C.3 Validating Shipping-Payment Dates with Retail Product Entry

The identifying idea is simple: if payments to delivery and transportation providers mark the physical arrival of goods, they should concentrate when a newly purchased product becomes available for resale. For each retailer–product pair, we identify the week in which the product first appears in the retailer's sales and trace the probability that the retailer makes a delivery or transportation payment over the six weeks before and after that week. We retain all newly purchased products, regardless of whether their first sale precedes or follows the shipping payment.

Retailers introduce new products repeatedly, and different product-entry events for the same retailer frequently overlap. A conventional event study centered on a single treatment date is therefore not well suited to this setting. We instead estimate separate horizon-specific local projections in the spirit of Jordà (2005); Jordà and Taylor (2025).

Figure C.6 shows a single, pronounced spike at $h = 0$, the week in which the product first appears in the retailer's sales. The increase in the probability of a delivery or transportation payment in that week is nearly 200% of the dependent-variable mean. The estimates at the six preceding and six subsequent horizons are much smaller. The concentration of the response precisely when the product enters downstream retail sales provides validation that shipping-payment dates proxy for physical product delivery.

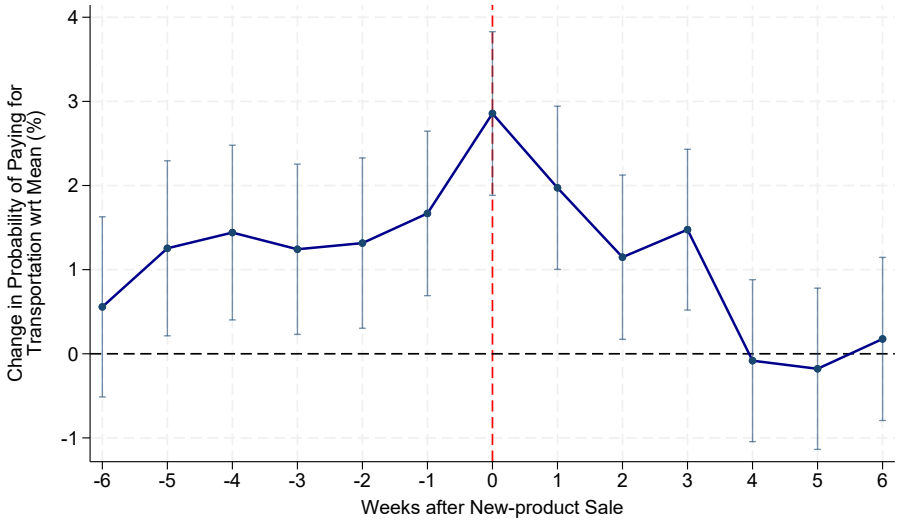


Figure C.6: Delivery and Transportation Payments Spike at Retail Product Entry

Notes: The figure reports horizon-specific local-projection estimates for $h = -6, \ldots, 6$. Time zero is the week in which a newly purchased product first appears in the retailer's sales. The dependent variable indicates whether the retailer makes a payment to a delivery or transportation provider. The sample includes all newly purchased products, regardless of whether their first sale precedes or follows the shipping payment. Separate horizon-specific regressions accommodate repeated and overlapping product-entry events within retailers. The magnitude discussed in the text expresses the $h = 0$ coefficient relative to the dependent-variable mean.

C.4 Price Premium in the Shipping-Observed Sample

Column (4) of Table 2 requires observing a firm's use of external delivery or transportation services. Table C.7 thus repeats all four specifications on the column-(4) estimation sample. The benefit of this restriction is that the sample remains fixed across columns; the cost is that firms never observed paying for external shipping are excluded.

Table C.7: Deferred Settlement, Shipping Timing, and Prices: Common Sample

Dependent variable: log transaction price				
	(1)	(2)	(3)	(4)
Deferred-payment invoice	0.226***	0.162***	0.162***	
	(0.005)	(0.005)	(0.005)	
Deferred payment: shipping at or after payment deadline				0.152***
				(0.005)
Deferred payment: shipping before payment deadline				0.170***
				(0.005)
Product fixed effects	Yes	No	No	No
Seller fixed effects	Yes	No	No	No
Seller × product fixed effects	No	Yes	Yes	Yes
Seller × date fixed effects	No	No	Yes	Yes
Product × date fixed effects	No	No	Yes	Yes
Calendar-date fixed effects	Yes	Yes	No	No
Adjusted R^2	0.674	0.775	0.787	0.787
Observations	74,534,406	74,534,406	74,534,406	74,534,406
Buyer clusters	298,493	298,493	298,493	298,493

Notes: The unit of observation is a transaction. The sample is held fixed across all columns and corresponds to the estimation sample used in column (4) of Table 2. It includes firms for which the timing of external delivery or transportation payments can be measured and excludes firms never observed purchasing such services. In columns (1)–(3), the independent variable is an indicator equal to one if payment is deferred rather than immediate. Column (1) includes product, seller, and exact calendar-date fixed effects. Column (2) includes seller-by-product and exact calendar-date fixed effects. Column (3) includes seller-by-product, seller-by-date, and product-by-date fixed effects. Column (4) retains the fixed effects in column (3) and divides deferred transactions according to whether the associated delivery-service payment is observed before the full-payment due date or at or after the full payment comes due. Immediate-payment transactions are the omitted category. Repayment windows longer than 360 days are excluded. Robust standard errors clustered at the buyer level are reported in parentheses. The data span 2024. *** denotes significance at the 1 percent level.

C.5 Timing-Based Chained Payments and Prices

Table C.8: Timing-Based Chained Payments and Prices

Dependent variable: Price of transaction t for good p sold by firm j to firm k at time t (asinh)

	(1)
$Chained_{kt} \times Deferred\,payment_{vpjt}$	0.106***
	(0.005)
$Deferred\,payment_{vpjt}$	0.031***
	(0.005)
$Chained_{kt}$	0.052***
	(0.001)
Observations	45,742,078
Adjusted R^2	0.785
$j \times p\,t$ FE	Yes

Notes: The regression includes an indicator for deferred payment, an indicator for whether buyer k is classified as chained, and their interaction. The unit of observation is a transaction. Seller×product and time fixed effects are included. Robust standard errors clustered at the buyer level are reported in parentheses. The data span 2024. *** denote significance at the 1 percent level.

C.6 Firm-Level TFP Estimation

We estimate monthly firm-level TFP using the Levinsohn–Petrin method (Levinsohn and Petrin, 2003). Firms are consolidated at the corporate-group level. Log sales are the output measure, employment is the freely variable input, capital is the state variable, and purchases net of capital expenditures are the proxy input. The capital stock starts from 2023 net assets and adds cumulative 2024 capital purchases. We estimate the production function and use the resulting residual as monthly revenue-based TFP.

The monthly frequency is deliberate: it allows us to align firm performance with the short-run timing of the Desyfin freeze and pre-existing contractual due dates, which annual data would obscure. We estimate a common production function over the corporate-group panel and apply the same residual construction to every firm-month. Sales, purchases, labor, and capital are aggregated consistently at the corporate-group-by-month level, with purchases net of capital expenditures serving as the Levinsohn–Petrin proxy variable. Because output is measured using invoiced sales, the resulting residual is a measure of monthly revenue productivity, the empirical counterpart of the model's measured TFP. We use this residual as one outcome in the reduced-form de-

sign; Desyfin exposure and contractual due-date proximity are constructed separately from the freeze and pre-existing invoice obligations. The same contractual-timing gradient also appears in sales and value added, which do not depend on the Levinsohn–Petrin inversion. Thus, the evidence that the payment disruption affects real activity does not rely on the productivity residual alone.

C.7 Earlier-Outstanding-Invoice Placebo

We test whether the main due-date gradient is driven by persistent characteristics of firms or their contracting arrangements. For each link, we reconstruct the contractual-timing measure using the deferred invoices outstanding on June 13, 2024, two months before the Desyfin intervention, applying the same invoice-selection rule used for the main measure.

Let $d_{i,J}^{(\ell)}$ denote the contractual due date selected from the invoices outstanding on June 13 for seller i at link ℓ, and define

$$D_{i,J}^{(\ell)} = d_{i,J}^{(\ell)} - \text{June 13, 2024}, \qquad \ell \in \{1, 2\}.$$

We replace the contractual-timing measure constructed from the invoices outstanding at the August freeze with $(D_{i,J}^{(\ell)})^{-1}$, while retaining the actual Desyfin-exposure indicator, the actual post-freeze period, and the remaining specification. Importantly, this measure is constructed from a different set of outstanding payment obligations; it is not obtained by applying an earlier reference date to the August invoices.

Table C.9: Earlier-Outstanding-Invoice Placebo

	TFP	Sales	Value Added
	(1)	(2)	(3)
Panel A. Sellers one link upstream			
$(D_{\mathrm{P}}^{(1)})^{-1} \times \mathrm{DesyfinLinked}^{(1)} \times \mathrm{Shock}$	-0.045	-0.015	-0.020
	(0.048)	(0.052)	(0.047)
Lower-order interactions	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Adjusted R^2	0.650	0.955	0.891
Observations	30,031	36,062	30,577
Clusters	685	748	740
Panel B. Sellers two links upstream			
$(D_{\mathrm{P}}^{(2)})^{-1} \times \mathrm{DesyfinLinked}^{(2)} \times \mathrm{Shock}$	0.003	0.014	0.004
	(0.012)	(0.019)	(0.021)
Lower-order interactions	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Adjusted R^2	0.596	0.934	0.852
Observations	64,054	92,533	75,143
Clusters	965	1,094	1,082

Notes: The table repeats the one- and two-link specifications from Table 4, replacing the contractual-timing measure constructed from invoices outstanding at the August 13 Desyfin freeze with one reconstructed from deferred invoices outstanding on June 13, 2024, and their own contractual due dates. The actual Desyfin-exposure indicators, actual post-freeze period, lower-order interactions, fixed effects, and clustering structure are otherwise unchanged. The placebo sample includes firms for which the corresponding June outstanding invoice can be constructed. Sales and value added enter in logs. TFP is estimated as described in Appendix C.6. Robust standard errors are clustered at the due-invoice buyer level.

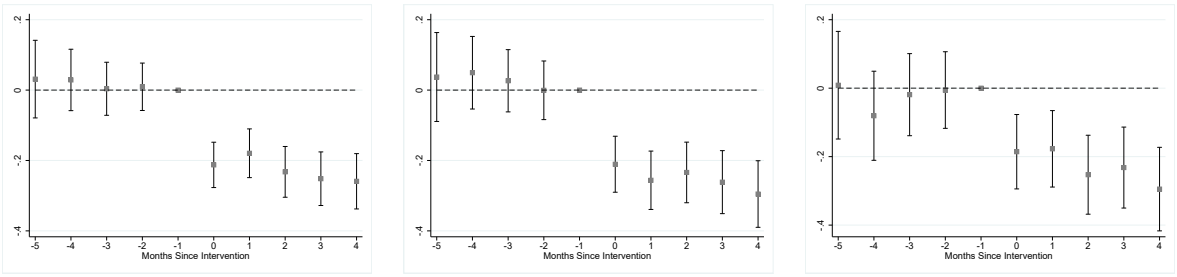
C.8 Complementary Timing-Based Evidence from the Desyfin Freeze

The invoice-based design in equation (5) identifies upstream real effects using payment obligations and contractual due dates fixed before the intervention. This appendix examines the accompanying adjustment in payment behavior. Column (1) of Table C.10 reports whether directly exposed firms become more likely to switch from spot to chained payments. Columns (2)–(7) describe activity changes among exposed firms and their upstream suppliers conditional on realized switching. Because switching can itself respond to the freeze and to changes in firm activity, we interpret these activity coefficients as conditional associations that characterize the joint adjustment of payments and real activity. The broader sample allows us to trace these patterns over time and across successive upstream links. Figure C.7 reports the dynamic response among suppliers one link upstream, while Figure C.8 traces how the effects vary with distance along the payment chain.

Table C.10: A Payments-Chains Crisis: Frozen Bank Accounts

A. Directly Shocked Firm i				
$Switch2chained_{it} \times Desyfin_i \times Shock_t$		-0.176^{***}	-0.233^{***}	-0.197^{*}
		(0.064)	(0.086)	(0.109)
$Switch2chained_{it} \times Shock_t$		-0.004	0.019^{*}	-0.017
		(0.577)	(0.010)	(0.012)
$Desyfin_i \times Shock_t$	0.016^{**}	0.159^{***}	0.195^{***}	0.204^{***}
	(0.007)	(0.048)	(0.037)	(0.048)
$\%\Delta$ w.r.t. mean [mean dep. var]	9% [0.18]			
Observations	1,200,404	1,192,436	939,984	636,485
Firm and time fixed effects	Yes	Yes	Yes	Yes
B. Firm j Selling to a Directly Shocked Firm i				
		ln $Purchases_{jt}$	ln $Sales_{jt}$	ln $Value\ Added_{jt}$
		(5)	(6)	(7)
$Selling2Desyfin_{jt} \times Switch2chained_{it}$		-0.176^{***}	-0.189^{***}	-0.130^{***}
$\times Desyfin_i \times Shock_t$		(0.029)	(0.033)	(0.038)
$Switch2chained_{it} \times Shock_t$		0.038^{**}	0.017	0.000
		(0.016)	(0.019)	(0.023)
$Selling2Desyfin_{jt} \times Switch2chained_{it}$		0.651^{***}	0.669^{***}	0.661^{***}
		(0.040)	(0.044)	(0.047)
$Desyfin_i \times Shock_t$		0.035^{**}	0.027	0.000
		(0.016)	(0.019)	(0.023)
Observations		704,087	555,098	382,319
Firm and time fixed effects	Yes	Yes	Yes	Yes

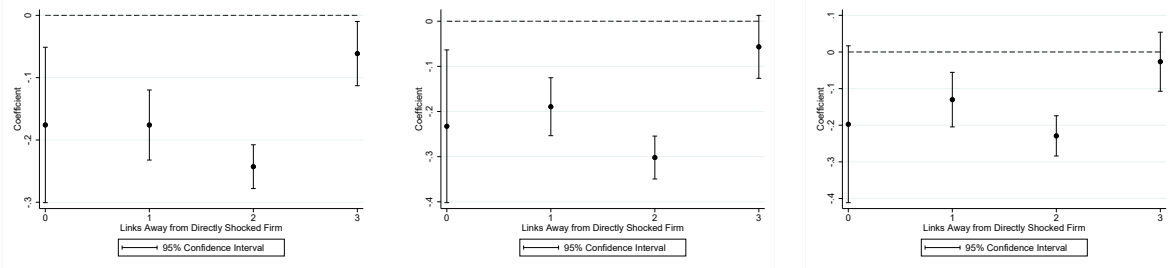
Notes: Column (1) examines switching as an outcome of exposure. Columns (2)–(7) report activity changes conditional on realized switching and describe the joint adjustment of payments and activity. The preferred estimates based on pre-intervention invoice obligations and due dates are reported in equation (5) and Table 4. The dependent variable in column (1) is an indicator equal to 1 if firm i switched from being classified as a spot firm to a chained firm in period t. Columns (2), (3), and (4) investigate the impact of the freeze on firms who held accounts at Desyfin, before and after the freeze, differentiating between impacts if the firm switched to being classified as chained or not. Columns (5), (6), and (7) instead focus on firms that were selling to a firm directly connected to Desyfin, before and after the shock. Robust standard errors, clustered by firm, are in parentheses. All specifications include firm and time fixed effects and lower-order terms implied by the interactions. Data are monthly and span 2024. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.



(a) $\ln Purchases_{jt}$ (b) $\ln Sales_{jt}$ (c) $\ln Value\ Added_{jt}$

Figure C.7: Upstream Activity Changes Around the Desyfin Freeze

Notes: The figure shows an event study where time zero corresponds with the month when Desyfin accounts were frozen. The dependent variables in panels (a), (b), and (c) are the (log) of purchases, of sales, and of value added, respectively, of firm j in month t. The plotted coefficients correspond with the triple interaction of $Selling2Desyfin_{jt} \times Switch2chained_{it} \times Desyfin_i \times Shock_t$, and therefore captures the impact of being one link upstream from a Desyfin firm *and* the differential activity changes associated with realized switching of firm i. Robust standard errors are clustered by firm. All specifications include firm and time fixed effects and the lower-order terms implied by the interactions. The data are monthly and span 2024.



(a) $\ln Purchases_{jt}$ (b) $\ln Sales_{jt}$ (c) $\ln Value\ Added_{jt}$

Figure C.8: Activity Changes by Distance from Directly Exposed Firms

Notes: The figure shows how the association with exposure varies with the number of links separating a firm from a directly exposed Desyfin firm. The dependent variables in panels (a), (b), and (c) are the (log) of purchases, sales, and value added, respectively, of firm j in month t. The plotted coefficients correspond with the triple interaction of $Links(n)_{jt} \times Switch2chained_{it} \times Desyfin_i \times Shock_t$, and capture the impact of being n links away from a Desyfin firm i ($Links(n)_{jt}$) *and* the differential activity changes associated with realized switching of firm i. Robust standard errors are clustered by firm. All specifications include firm and time fixed effects and the lower-order terms implied by the interactions. The data are monthly and span 2024.

D Proofs for the Partial Equilibrium Analysis (Section 4)

D.1 Entrepreneur's Problem: Details

This appendix provides a detailed characterization of the entrepreneur's problem, including formal statements and proofs of the key results presented in the main text.

Thresholds. The following lemma reproduces Lemma 1 in the main text, which identifies two debt threshold points that determine whether an entrepreneur's expenditure is a mix between spot and chained purchases or whether expenditures are concentrated in one type of order.

Lemma 10. (Expenditure Threshold Points): *Define the* ***efficiency threshold****,* $B^{\star}_{t+1} \equiv R_{t+1}\left(\tilde{B}_t - 1\right)$*. Then:*

(i) $S_{t+1} = 0$ *if and only if* $B_{t+1} \geq \tilde{B}_{t+1}$.

(ii) If $B_{t+1} < \tilde{B}_{t+1}$, $X_t > 0$ *if and only if* $B_{t+1} > B^{\star}_{t+1}$.

Proof. Part (i): If $B_{t+1} \geq \tilde{B}_{t+1}$, then $\bar{S}_{t+1} = \max\{\tilde{B}_{t+1} - B_{t+1}, 0\} = 0$. By conditions (6), $S_{t+1} = \min\{\bar{S}_{t+1}, E_{t+1}\} = 0$. Conversely, if $S_{t+1} = 0$, then either $\bar{S}_{t+1} = 0$ or $E_{t+1} = 0$. Since entrepreneurs must have positive endowments at all t consume at all periods, $E_{t+1} > 0$, implying $\bar{S}_{t+1} = 0$, which occurs if and only if $B_{t+1} \geq \tilde{B}_{t+1}$.

Part (ii): From equation (6), chained expenditures are positive, $X_t > 0$, if and only if expenditures exceed the limit on spot expenditures, $E_t > \bar{S}_t = \max\{\tilde{B}_t - B_t\, 0\}$. Using the budget constraint $E_t = B_{t+1}R^{-1}_{t+1} + 1 - B_t$, this condition becomes:

$$B_{t+1}R^{-1}_{t+1} + 1 - B_t > \max\{\tilde{B}_t - B_t, 0\}.$$

Rearranging terms yields $B_{t+1} > R_{t+1}(\max\{\tilde{B}_t, B_t\} - 1) = B^{\star}_{t+1}(\tilde{B}_t, B_t)$. □

As stated in the text, the efficiency threshold $B^{\star}_{t+1}$ marks the debt level above which entrepreneurs must resort to chained purchases in the current period due to insufficient spot borrowing capacity. The threshold $\tilde{B}_{t+1}$ determines whether entrepreneurs lose access to spot purchases entirely in the following period.

Marginal Prices and Optimality Conditions. The presence of two types of purchases—spot and chained—with different prices creates discontinuities in the marginal cost of buying goods. We formalize this through the concept of marginal prices, reproducing the definition in the main text.

Definition 3. (Average and Marginal Prices):

I. *The **average price** per unit of consumption is* $Q_t \equiv E_t/(S_t + X_t)$.

II. *The **marginal expenditure price** at* t *given debt choice* B' *is:*

$$\tilde{q}_t^E(B') \equiv 1 + (q_t - 1)\,\mathbb{I}_{\left[B' \geq B_{t+1}^\star\right]}$$

III. *The **marginal borrowing price** at* $t+1$ *given debt* B' *is:*

$$\tilde{q}_{t+1}^B(B') \equiv 1 + (q_{t+1} - 1)\,\mathbb{I}_{\left[B' > \tilde{B}_{t+1}\right]}$$

The marginal expenditure price $\tilde{q}_t^E(B')$ represents the price at which the entrepreneur purchases goods at the margin in period t. If the choice of future debt exceeds the efficiency threshold, $B' \geq B_{t+1}^\star$, the marginal price of goods equals q_t, the value of chained goods; otherwise, it equals 1 (the spot price). Similarly, $\tilde{q}_{t+1}^B(B')$ captures the price at which saved resources translate into consumption in period $t+1$. If future debt exceeds the SBL ($B' > \tilde{B}_{t+1}$), any savings yield chained purchases at price q_{t+1} because no spot goods can be purchased; otherwise, they yield spot purchases at price 1 because any additional unit of savings allows an additional unit of spot consumption on the margin.

D.2 Proof of Proposition 2 (Entrepreneur's Euler Equation)

We begin by formulating the entrepreneur's problem in sequential form:

Problem 5. *Given* B_0 *and* $\left\{R_{t+1}, \tilde{B}_t\right\}_{t\geq 0}$,

$$\max_{\{S_t, X_t\}_{t\geq 0}} \sum_{t\geq 0} \beta^t \log\left(C_t\right),$$

subject to the budget constraint, $B_t + E_t = R_{t+1}^{-1} B_{t+1} + 1, \forall t \geq 0$, *conditions* (6) *and* $C_t = S_t + X_t$, *and to the intra- and inter-period constraints,* $S_t \leq \bar{S}_t$ *and* $B_t \leq \bar{B}$.

We derive the entrepreneur's Euler equation as a necessary condition for optimality. The entrepreneur's total expenditures as a function of B_{t+1} are:

$$E_t = E(B_t, B_{t+1}, R_{t+1}) \equiv 1 - B_t + \frac{B_{t+1}}{R_{t+1}}.$$

Given total expenditures, spot expenditures are:

$$S\left(B_t, \tilde{B}_t, B_{t+1}\right) = \min\left\{\max\left\{\tilde{B}_t - B_t, 0\right\}, 1 - B_t + \frac{B_{t+1}}{R_{t+1}}\right\},$$

and chained expenditures are:

$$q_t X_t = E(B_t, B_{t+1}, R_{t+1}) - S\left(B_t, \tilde{B}_t, B_{t+1}\right).$$

The entrepreneur's consumption (total goods obtained) is:

$$C_t = S_t + X_t \quad = \quad S\left(B_t, \tilde{B}_t, B_{t+1}\right) + \frac{E(B_t, B_{t+1}, R_{t+1}) - S\left(B_t, \tilde{B}_t, B_{t+1}\right)}{q_t}.$$

Thus, we can write current consumption as:

$$C_t = C(B_t, B_{t+1}, \tilde{B}_t, q_t, R_{t+1}) \equiv \frac{E(B_t, B_{t+1}, R_{t+1})}{q_t} + \left(1 - \frac{1}{q_t}\right) S\left(B_t, \tilde{B}_t, B_{t+1}\right).$$

The entrepreneur's intertemporal optimization problem can be written entirely in terms of a sequence of debt levels:

$$\max_{\{B_{t+1}\}_{t\geq 0}} \sum_{t\geq 0} \beta^t \log\left(C(B_t, B_{t+1}, \tilde{B}_t, q_t, R_{t+1})\right)$$

with B_0 given.

Consider the choice of B_{t+1}. Define:

$$\Upsilon_t\left(B_{t+1};B_t,B_{t+2},\tilde{B}_t,\tilde{B}_{t+1},R_{t+1},R_{t+2},q_t,q_{t+1}\right) \equiv \log\left(C(B_t,B_{t+1},\tilde{B}_t,q_t,R_{t+1})\right) + \beta\log\left(C(B_{t+1},B_{t+2},\tilde{B}_{t+1},q_{t+1},R_{t+2})\right).$$

The function Υ_t is the component of the objective that contains the influence of B_{t+1}. Then, we have that Υ_t features kinks at the threshold points $\left\{B^\star_{t+1},\tilde{B}_{t+1}\right\}$. These kinks arise from the kinks in the functions E and S–reflecting forced switches between spot and chained purchases.

We analyze the optimality conditions by studying the effect of a marginal change in B_{t+1}. Recall the relation between $\tilde{B}_t$ and $B^\star_{t+1}$:

$$B^\star_{t+1} = R_{t+1}\left(\tilde{B}_t - 1\right).$$

We do so by splitting the marginal conditions depending on whether some spot expenditures are available at t:

Case I: $B_t \geq \tilde{B}_t$ (No spot purchases at t). We distinguish subcases depending on whether debt levels permit spot expenditures tomorrow:

I.a. If $B_{t+1} < \tilde{B}_{t+1}$ (spot purchases available at $t+1$):

$$\Upsilon_t'(B_{t+1}) = \frac{1}{C_t}\frac{1}{q_t R_{t+1}} - \beta\frac{1}{C_{t+1}}.$$

I.b. If $B_{t+1} > \tilde{B}_{t+1}$ (no spot purchases at $t+1$):

$$\Upsilon_t'(B_{t+1}) = \frac{1}{C_t}\frac{1}{q_t R_{t+1}} - \beta\frac{1}{C_{t+1}}\frac{1}{q_{t+1}}$$

At $B_{t+1} = \tilde{B}_{t+1}$, the derivative features a discontinuity:

$$\lim_{B'\uparrow\tilde{B}_{t+1}}\Upsilon_t'(B') < \lim_{B'\downarrow\tilde{B}_{t+1}}\Upsilon_t'(B'),$$

where the inequality follows from q_{t+1}. The condition implies that the marginal increase in utility is greater to the right $\tilde{B}_{t+1}$. Since a solution at the kink would require:

$$\lim_{B'\uparrow\tilde{B}_{t+1}}\Upsilon_t'(B') > 0 > \lim_{B'\downarrow\tilde{B}_{t+1}}\Upsilon_t'(B'),$$

no solution will feature $B_{t+1} = \tilde{B}_{t+1}$.

Case II: $B_t < \tilde{B}_t$ (Possible spot purchases at t). We distinguish subcases based on the debt level relative to the thresholds:

II.a. If $B_{t+1} < B^\star_{t+1}$ (only spot at t, some spot at $t+1$):

$$\Upsilon'_t(B_{t+1}) = \frac{1}{C_t} \cdot \frac{1}{R_{t+1}} - \beta \frac{1}{C_{t+1}}$$

II.b. If $B_{t+1} \in \left(B^\star_{t+1}, \tilde{B}_{t+1}\right)$ (mixed at t, spot at $t+1$):

$$\Upsilon'_t(B_{t+1}) = \frac{1}{q_t}\frac{1}{C_t} \cdot \frac{1}{R_{t+1}} - \beta \frac{1}{C_{t+1}}$$

II.c. If $B_{t+1} > \tilde{B}_{t+1}$ (mixed at t, only chained at $t+1$):

$$\Upsilon'_t(B_{t+1}) = \frac{1}{C_t} \cdot \frac{1}{q_t}\frac{1}{R_{t+1}} - \beta \frac{1}{C_{t+1}}\frac{1}{q_{t+1}}$$

At the threshold points, we have the following discontinuities:

$$\lim_{B' \uparrow B^\star_{t+1}} \Upsilon'_t(B') > \lim_{B' \downarrow B^\star_{t+1}} \Upsilon'_t(B'), \quad \lim_{B' \uparrow \tilde{B}_{t+1}} \Upsilon'_t(B') < \lim_{B' \downarrow \tilde{B}_{t+1}} \Upsilon'_t(B') \tag{75}$$

Using the definition of average price $Q_t = E_t/C_t$ and marginal prices, we express Υ'_t in a compact form:

$$\Upsilon'_t(B_{t+1}) = \frac{Q_t}{E_t} \cdot \frac{1}{R_{t+1}\tilde{q}^E_t(B_{t+1})} - \beta \frac{Q_{t+1}}{E_{t+1}} \cdot \frac{1}{\tilde{q}^B_{t+1}(B_{t+1})}$$

Setting $\Upsilon'_t(B_{t+1}) = 0$ at interior solutions yields:

$$\frac{E_{t+1}/E_t}{\beta} \cdot \frac{Q_t}{Q_{t+1}} = \frac{R_{t+1}}{\tilde{q}^B_{t+1}(B_{t+1})/\tilde{q}^E_t(B_{t+1})}. \tag{76}$$

We already established that $B_{t+1} = \tilde{B}_{t+1}$ is not a possible outcome. Hence, the only relevant corner solution is $B_{t+1} = B^\star_{t+1}$. Optimality at that corner solution is given by:

$$\lim_{B' \uparrow B^\star_{t+1}} \Upsilon'_t(B') \geq 0 \geq \lim_{B' \downarrow B^\star_{t+1}} \Upsilon'_t(B'),$$

Opening up the condition, using that average prices are $Q_t = Q_{t+1} = 1$ at a corner translates into:

$$\beta q_t R_{t+1} \geq \frac{E_{t+1}}{E_t} \geq \beta R_{t+1}. \tag{77}$$

We arrive at equations (77) and (76), which completes the proof of Proposition 2. □

The Euler equation captures how the entrepreneur's consumption-saving decision depends on both current and future marginal prices, which vary discontinuously with debt levels due to the spot borrowing line constraint.

D.3 Proof of Proposition 3 (Entrepreneur's Stationary Solution)

In the stationary version of the entrepreneur's problem, $R_t = \beta^{-1}$, the SBL is constant, $\tilde{B}_t = \tilde{B}_{ss}$, and the chained-goods price is constant, $q_t = q_{ss}$. Throughout this proof, we assume $\tilde{B}_{ss} > 1$. For convenience, we suppress the arguments of $B^\star(1/\beta, \tilde{B}_{ss})$ and denote its steady-state value simply as $B^\star_{ss}$.

We establish the following results:

I. If $B_0 \in [0, B^\star_{ss}]$, then $B_t = B_0$ for all t.

II. There exists a threshold $B^h > \tilde{B}_{ss}$ such that:

II.a If $B^\star_{ss} < B_0 < B^h$, then $B_t \to B^\star_{ss}$ in finite time.

II.b If $B_0 > B^h$, then $B_t = B_0$ for all t.

The proof strategy involves: First, establishing two benchmark value functions, $\bar{V}(B)$ and $\underline{V}(B)$, that bound the entrepreneur's value function. Second, showing that for $B_0 \leq B^\star_{ss}$, a constant debt policy is optimal and achieves the upper bound $\bar{V}(B)$. Third, proving that for $B_0 \in (B^\star_{ss}, \tilde{B}_{ss}]$, optimal deleveraging occurs in finite time with debt converging to $B^\star_{ss}$. Finally, establishing the existence and uniqueness of $B^h > \tilde{B}_{ss}$ such that entrepreneurs with debt above B^h maintain constant debt and $V(B) = \underline{V}(B)$.

Value Function Bounds. Consider two auxiliary problems that provide bounds for the entrepreneur's value function. First, the unconstrained entrepreneur problem:

Problem 6 (Upper Bound)**.**

$$\bar{V}(B) = \max_{B' \leq \bar{B}} \ln(E) + \beta \bar{V}(B')$$

subject to: $B + E = 1 + \frac{B'}{R}$.

Second, an entrepreneur problem where only chained goods are consumed:

Problem 7 (Lower Bound)**.**

$$\underline{V}(B) = \max_{B' \leq \bar{B}} \ln(E/q_{ss}) + \beta \underline{V}(B')$$

subject to: $B + E = 1 + \frac{B'}{R}$.

Lemma 11. *The solutions to the benchmark problems are:*

$$\bar{V}(B) = \frac{\ln(1-(1-\beta)B)}{1-\beta}$$

and

$$\underline{V}(B) = \frac{\ln(1-(1-\beta)B) - \ln(q_{ss})}{1-\beta}.$$

In both cases, optimal expenditures are $E = 1-(1-\beta)B$.

Proof. For $\underline{V}$, we guess and verify $\underline{V}(B) = \frac{\ln(1-(1-\beta)B)-\ln(q_{ss})}{1-\beta}$. Substituting into the Bellman equation and taking first-order conditions with respect to B' yields:

$$\frac{1}{1-B+B'/R}\frac{1}{R} = \frac{\beta}{1-(1-\beta)B'}.$$

Using $\beta = R^{-1}$, we obtain $B = B'$, confirming constant debt. Expenditures are $E = 1-(1-\beta)B$, verifying the conjecture. Setting $q_{ss} = 1$ yields $\bar{V}(B)$. □

Note that $V(B) \in [\underline{V}(B), \bar{V}(B)]$ since these correspond to more constrained and relaxed versions of the original problem.

Proof of Part I.

Lemma 12. *For* $B \leq B^{\star}_{ss}$, $V(B) = \bar{V}(B)$ *and* $E = 1-(1-\beta)B$.

Proof. We verify that setting $E = 1-(1-\beta)B$ in the original problem yields the upper bound value. Note that:

$$\begin{aligned} E = 1-(1-\beta)B \leq \tilde{B}_{ss} - B \quad &\iff \quad 1+\beta B \leq \tilde{B}_{ss} \\ &\iff \quad B \leq \beta^{-1}(\tilde{B}_{ss}-1) = B^{\star}_{ss}. \end{aligned}$$

Thus, consuming only through spot purchases is feasible. With $B' = \beta^{-1}(B+E-1) = B$ and consumption $C = S_t + X_t = 1-(1-\beta)B$ constant across periods, $V(B)$ attains the upper bound. □

Proof of Part II - Preliminary Results. Define the set of debt levels from which setting $B' = B^\star_{ss}$ is optimal:

$$\mathcal{B}_0 \equiv \left[B^\star_{ss}, B^\star_{ss} + (1 - \frac{1}{q_{ss}})(\tilde{B}_{ss} - B^\star_{ss}) \right]. \tag{78}$$

Lemma 13. *If $B' = B^\star_{ss}$, then $C = \tilde{B}_{ss} - B$.*

Proof. When $B' = B^\star_{ss}$, by definition $Q = 1$. Thus:

$$C = 1 + \frac{B^\star_{ss}}{R} - B = 1 + \frac{(\tilde{B}_{ss} - 1)}{\beta R} - B = \tilde{B}_{ss} - B,$$

using $B^\star_{ss} = 1/\beta(\tilde{B}_{ss} - 1)$. □

Proposition 6. *$B' = B^\star_{ss}$ if and only if $B \in \mathcal{B}_0$.*

Proof. We find values of B such that the sub-differential condition holds:

$$\frac{C(B^\star_{ss})}{C(B)} \in [1, q_{ss}],$$

where marginal prices are both 1 since $B' \leq B^\star_{ss} < \tilde{B}_{ss}$ and $\beta R = 1$. This yields:

$$\frac{\tilde{B}_{ss} - B^\star_{ss}}{\tilde{B}_{ss} - B} \geq 1 \iff B \geq B^\star_{ss},$$

and

$$\frac{\tilde{B}_{ss} - B^\star_{ss}}{\tilde{B}_{ss} - B} \leq q_{ss} \iff B \leq \tilde{B}_{ss} - \frac{1}{q_{ss}}(\tilde{B}_{ss} - B^\star_{ss}) < \tilde{B}_{ss}.$$

□

Proposition 7. *Let $\gamma \equiv \beta/q_{ss}$ and define:*

$$\Gamma(t) \equiv \frac{1 - \gamma^t}{1 - \gamma} - q_{ss}^{-t} \frac{1 - \beta^t}{1 - \beta}, \quad t \geq 0.$$

Let T be the unique positive integer satisfying $\Gamma(T) < 1 < \Gamma(T + 1)$.

Any decreasing sequence $\{B_{-t}\}$ for $t \in \{0, T\}$ satisfying the entrepreneur's Euler equation with terminal condition $B_0 \in \mathcal{B}_0$ satisfies $B_{-t} \in \mathcal{B}_{-t}$ where

$$\mathcal{B}_{-t} \equiv (B^\ell_{-t}, B^\ell_{-t-1}] \cap [0, \tilde{B}_{ss}],$$

with

$$B^\ell_{-t} \equiv B^\star_{ss} + (\tilde{B}_{ss} - B^\star_{ss})\Gamma(t).$$

Moreover, $\cup_{t=0}^{T} \mathcal{B}_{-t} = (B^{\star}_{ss}, \tilde{B}_{ss}]$.

Proof. For any $B \in (B^{\star}_{ss}, \tilde{B}_{ss}]$, the entrepreneur makes both spot and chained purchases, so the ratio of marginal prices is q_{ss}. Consider a decreasing sequence of debt levels converging to some $B_0 \in \mathcal{B}_0$. For any $B_{-t} \in (B^{\star}_{ss}, \tilde{B}_{ss}]$, the entrepreneur's Euler equation yields:

$$\frac{C_{-t+1}}{C_{-t}} = \beta R \cdot \frac{\tilde{q}^{E}_{-t}(B_{-t+1})}{\tilde{q}^{B}_{-t+1}(B_{-t+1})} = q_{ss},$$

where we used that $\beta R = 1$ in steady state and the ratio of marginal prices equals q_{ss} for debt in this range. This gives us the difference equation:

$$C_{-t} = \frac{1}{q_{ss}} C_{-t+1} \text{ for } B_{-t+1} \in (B^{\star}_{ss}, \tilde{B}_{ss}],$$

with terminal condition $C_0 = \tilde{B}_{ss} - B_0$ (from the previous lemma).

Shifting forward the equation, we obtain:

$$C_0 / C_{-t} = q^{t}_{ss}.$$

Using the budget constraint, consumption at $-t$ is:

$$C_{-t} = \underbrace{\tilde{B}_{ss} - B_{-t}}_{S_{-t}} + \underbrace{\frac{1}{q_{ss}}(1 + \beta B_{-t+1} - \tilde{B}_{ss})}_{X_{-t}}.$$

Combining these expressions:

$$\tilde{B}_{ss} - B_{-t} + q^{-1}_{ss}(1 + \beta B_{-t+1} - \tilde{B}_{ss}) = q^{-t}_{ss}(\tilde{B}_{ss} - B_0).$$

Rearranging:

$$B_{-t} = \tilde{B}_{ss} - q^{-t}_{ss}(\tilde{B}_{ss} - B_0) + q^{-1}_{ss}(1 + \beta B_{-t+1} - \tilde{B}_{ss}).$$

Since $1 + \beta B_{-t+1} - \tilde{B}_{ss} = \beta(B_{-t+1} - B^{\star}_{ss})$ (using the definition of $B^{\star}_{ss}$), we get:

$$B_{-t} = \tilde{B}_{ss} - q^{-t}_{ss}(\tilde{B}_{ss} - B_0) + \frac{\beta}{q_{ss}}(B_{-t+1} - B^{\star}_{ss}).$$

For the boundary sequence with $B_0 = B_0^\ell = B_{ss}^\star$, solving this difference equation yields:

$$B_{-t}^\ell = B_{ss}^\star + (\tilde{B}_{ss} - B_{ss}^\star) \sum_{\tau=0}^{t-1} \gamma^\tau (1 - q_{ss}^{-t+\tau}),$$

where $\gamma = \beta/q_{ss}$. This sum simplifies to $\Gamma(t)$ as defined.

Since $B_{-1}^\ell = B_{ss}^\star + (1 - 1/q_{ss})(\tilde{B}_{ss} - B_{ss}^\star) = B_0^u$ (the upper bound of $\mathcal{B}_0$), the intervals $\mathcal{B}_{-t}$ form a partition of $(B_{ss}^\star, \tilde{B}_{ss}]$. The condition $\Gamma(T) < 1 < \Gamma(T+1)$ ensures that T is the last period for which $B_{-T}^\ell \leq \tilde{B}_{ss}$. Since $\Gamma(1) = 1 - q_{ss}^{-1} < 1$ and $\lim_{t\to\infty} \Gamma(t) = 1/(1-\gamma) > 1$, with Γ increasing in t, there exists a unique such T. □

Existence and Uniqueness of B^h. Following the analysis for sequences above $\tilde{B}_{ss}$, we establish:

Corollary 3. *There exists a maximal debt level*

$$\underline{B} = \frac{1 - q_{ss}^{-T}(1 - (1-\beta)B_{ss}^\star)}{1-\beta},$$

such that deleveraging toward $B_{ss}^\star$ following the Euler equation is infeasible.

For values above $\underline{B}$, the policy $B' = B$ is the unique solution consistent with the Euler equation. Since $V(\tilde{B}_{ss}) > \underline{V}(\tilde{B}_{ss})$ while $V(\underline{B}) = \underline{V}(\underline{B})$, by continuity there exists $B^h \in [\tilde{B}_{ss}, \underline{B}]$ such that $V(B) = \underline{V}(B)$ for all $B > B^h$.

To establish uniqueness, we show that if $V(B^d) > \underline{V}(B^d)$ for some $B^d > \tilde{B}_{ss}$, then $V(B) > \underline{V}(B)$ for all $B < B^d$. This follows from the decreasing absolute risk aversion property of log utility: starting from any $B < B^d$, the entrepreneur can consume the annuity of the difference $\Delta = B^d - B$ in chained goods and follow the same deleveraging path, yielding strictly higher value than $\underline{V}(B)$.

This completes the proof of Proposition 3. □

E Proofs in the General Equilibrium Analysis (Section 5)

E.1 Proof of Proposition 4

Proof. The proof uses the individual optimization creditors together with the market-clearing conditions to arrive at a mapping from current debt levels to allocations:

Step 1. From the creditor's log utility and budget constraint $B_{t+1}/R_{t+1} + C_t = B_t$, optimal consumption is:

$$C(B) = (1-\beta)B \tag{79}$$

This is the annuity value of debt holdings.

Step 2. Market clearing for goods requires total expenditures equal total production (normalized to 1):

$$C(B) + E(B) = 1$$

Hence, entrepreneur expenditures are given by the function:

$$E(B) \equiv 1 - C(B) = 1 - (1-\beta)B. \tag{80}$$

Step 3. Given total expenditures $E(B)$ and the constraint that $q \geq 1$ (with equality only when $\mu = 0$), the entrepreneur maximizes goods obtained by maximizing spot purchases subject to the SBL:

$$S(B, \tilde{B}) \equiv \min\{\bar{S}, E(B)\} = \min\{\max\{0, \tilde{B} - B\}, 1 - (1-\beta)B\}$$

Step 4. The value of chained expenditures follows as the residual:

$$\mu(B, \tilde{B}) \equiv E(B) - S(B, \tilde{B}) = 1 - (1-\beta)B - \min\{\max\{0, \tilde{B} - B\}, 1 - (1-\beta)B\}.$$

This yields:

$$\mu(B, \tilde{B}) = \begin{cases} 0 & B < B^\star(\tilde{B}), \\ 1 + \beta B - \tilde{B} & B \in [B^\star(\tilde{B}), \tilde{B}], \\ 1 - (1-\beta)B & B > \tilde{B}. \end{cases}$$

Step 5. The equilibrium price of chained goods is determined by the network technology:

$$q(B, \tilde{B}) = \mathcal{A}^{-1}(\mu(B, \tilde{B}))$$

where $\mathcal{A}(\cdot)$ is the network efficiency function.

Step 6. The quantity of chained goods obtained by dividing chained expenditures by their price:

$$X(B, \tilde{B}) = \frac{\mu(B, \tilde{B})}{q(B, \tilde{B})} = \mu(B, \tilde{B})\mathcal{A}(\mu(B, \tilde{B})).$$

Step 7. The average price per unit of consumption:

$$Q(B, \tilde{B}) = \frac{E(B)}{S(B, \tilde{B}) + X(B, \tilde{B})} = \frac{S(B, \tilde{B}) + q(B, \tilde{B})X(B, \tilde{B})}{S(B, \tilde{B}) + X(B, \tilde{B})}.$$

□

E.2 Analysis: Preliminary Observations

Marginal Expenditure Price as a function of B. The marginal expenditure price depends on both debt and the SBL: $\tilde{q}^E_t = \tilde{q}^E(B_t, \tilde{B}_t)$ where:

$$\tilde{q}^E(B, \tilde{B}) \equiv q(\mu(B, \tilde{B}))\mathbb{I}_{[B \geq B^\star(\tilde{B})]} + (1 - \mathbb{I}_{[B \geq B^\star(\tilde{B})]}).$$

Recall that $\mu(B, \tilde{B})$ is continuous, starts at zero, increases up to $B = \tilde{B}$, then decreases. Since q_t is monotone in μ_t, the marginal expenditure price follows the same pattern.

At the efficiency threshold:

$$\lim_{B \downarrow B^\star(\tilde{B})} \tilde{q}^E(B, \tilde{B}) = \lim_{\mu \downarrow 0} q(\mu) = \delta^{-1},$$

where δ is the network efficiency parameter. Thus $\tilde{q}^E$ has a discontinuity at $B^\star$ since $\lim_{B \uparrow B^\star(\tilde{B})} \tilde{q}^E = 1 \neq \delta^{-1}$.

Average-to-Marginal Price Ratio as function of B. The ratio $Q/\tilde{q}^E$ is crucial for equilibrium dynamics. For $B < B^\star(\tilde{B})$, since $q = 1$, we have $Q/\tilde{q}^E = 1$. For $B \geq \tilde{B}$, also $\tilde{q}^E = Q = q$, so $Q/\tilde{q}^E = 1$.

In the intermediate range $B \in [B^\star(\tilde{B}), \tilde{B}]$:

$$\frac{Q}{\tilde{q}^E} = \frac{1}{q} \cdot \frac{q}{(1 - S/E) + q(S/E)} = \frac{1}{(1 - S/E) + q(\mu(B, \tilde{B}))(S/E)}.$$

At $B = \tilde{B}$, we have $S/E = 0$, so the function is continuous there. However, as $B \downarrow B^\star(\tilde{B})$:

$$\frac{S}{E} \downarrow 1 \text{ and } q \downarrow \delta^{-1} > 1,$$

yielding $Q/\tilde{q}^E = \delta$ at the right limit of $B^\star$ and $Q/\tilde{q}^E = 1$ at the left limit—a discontinuity.

Analysis of the Marginal Borrowing Price. The marginal borrowing price $\tilde{q}^B_{t+1} = \tilde{q}^B(B_{t+1}, \tilde{B}_{t+1})$ where:

$$\tilde{q}^B(B, \tilde{B}) \equiv q(\mu(B, \tilde{B}))\mathbb{I}_{[B \geq \tilde{B}]} + (1 - \mathbb{I}_{[B \geq \tilde{B}]}).$$

Considering the discontinuities in μ:

$$\tilde{q}^B(B, \tilde{B}) = \begin{cases} 1 & B < B^\star(\tilde{B}), \\ 1 & B \in [B^\star(\tilde{B}), \tilde{B}], \\ q & B > \tilde{B}. \end{cases}$$

At $B = \tilde{B}$, the function has a discontinuity:

$$\lim_{B \downarrow \tilde{B}} \tilde{q}^B(B, \tilde{B}) = q(\mu(\tilde{B}, \tilde{B})) = q(1 - (1 - \beta)\tilde{B}) > 1.$$

The ratio $Q/\tilde{q}^B$:

$$\frac{Q}{\tilde{q}^B} = \begin{cases} 1 & B < B^\star(\tilde{B}), \\ Q & B \in [B^\star(\tilde{B}), \tilde{B}], \\ 1 & B > \tilde{B}, \end{cases}$$

where $\lim_{B \downarrow B^\star(\tilde{B})} Q = 1$.

Average Price Elasticity. The elasticity of average price with respect to debt in general equilibrium (after imposing $E = 1 - (1 - \beta)B$) is critical for policy analysis.

Starting from:

$$Q = \frac{1}{\frac{1}{q}(1 - S/E) + S/E},$$

we obtain:

$$\frac{\partial Q}{\partial B} = -Q \cdot \frac{(1 - S/E)\frac{\partial}{\partial B}[1/q] - (1/q - 1)\frac{\partial}{\partial B}[S/E]}{(1/q(1 - S/E) + S/E)}.$$

The derivative of the expenditure share is given by:

$$\frac{\partial}{\partial B}\left[\frac{S}{E}\right] = \begin{cases} 0 & B < B^\star(\tilde{B}) \\ \left[\frac{\tilde{B}-B}{1-(1-\beta)B}\right]\left(\frac{(1-\beta)}{1-(1-\beta)B} - \frac{1}{\tilde{B}-B}\right) & B \in [B^\star(\tilde{B}), \tilde{B}) \\ 0 & B \geq \tilde{B}. \end{cases}$$

This derivative is negative in the intermediate region and has discontinuities at the boundaries:

$$\frac{\partial}{\partial B}\left[\frac{S}{E}\right]\Big|_{B\uparrow\tilde{B}} = -\frac{1}{1-(1-\beta)\tilde{B}} < -1,$$

and

$$\frac{\partial}{\partial B}\left[\frac{S}{E}\right]\Big|_{B\downarrow B^\star} = \frac{(1-\beta)\tilde{B}_t - 1}{\Big((1-\beta)B^\star(\tilde{B}_t, B_t) - 1\Big)^2} < 0.$$

For the price derivative:

$$\frac{\partial}{\partial B}[q] = \frac{\partial q}{\partial \mu} \cdot \frac{\partial \mu}{\partial B} = q_\mu \mu \frac{\frac{\partial E}{\partial B} - \frac{\partial S}{\partial B}}{E},$$

where ϵ_μ^q denotes the elasticity of q with respect to μ. These elasticities determine the general equilibrium feedback effects between debt accumulation and payment chain formation.

E.3 Derivation of the map $\mathcal{B}$

For convenience, we reproduce the aggregate Euler equation (7) here:

$$\underbrace{\frac{B}{1-(1-\beta)B} \cdot Q\left(B, \tilde{B}\right)}_{\equiv \mathcal{E}\left(B;\tilde{B}\right)} = \underbrace{\frac{B'}{1-(1-\beta)B'} \frac{Q\left(B', \tilde{B}'\right)}{\Pi\left(B'; B, \tilde{B}, \tilde{B}'\right)}}_{\equiv \mathcal{E}'\left(B';B,\tilde{B},\tilde{B}'\right)}.$$

First, we show that for a subset of $B > \tilde{B}'$ there are two roots B' that solve $\mathcal{E} = \mathcal{E}'$. The roots satisfy $B_1' < \tilde{B}' < B_2'$ and $B_2' = B$ thus $R = \beta^{-1}$ and $q = q'$. We then show for any $B \in \left(\tilde{B}', B^h\right)$ the larger root B_2' cannot be an individual optimum during a transition.

The first result is summarized by the following Lemma.

Lemma 14. *There exists a threshold B^* such that for any $B \in [\tilde{B}', B^*]$, the equation $\mathcal{E} = \mathcal{E}'$ has two roots B': one root is $B = B'$ and the other satisfies $B' \leq \tilde{B}'$. For any $B > B^*$, only $B = B'$ is a solution. B^* solves:*

$$\mathcal{E}'\left(B^*; \tilde{B}, \tilde{B}, \tilde{B}'\right) = \lim_{B' \uparrow \tilde{B}'} \mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right).$$

The interpretation of B^* is obtained from the left panel of Figure 9: B^* is the highest value for which the Euler equation has a solution $B' \in (\tilde{B}', B^*)$.

Proof. Since $\tilde{B}' > \tilde{B}$, if $B' = B > \tilde{B}'$, the worker only spends chained expenditures. Thus,

$$\Pi\left(B'; B, \tilde{B}, \tilde{B}'\right) = 1 \text{ and } Q\left(B, \tilde{B}\right) = Q\left(B', \tilde{B}'\right) = q\left(1 - (1-\beta) B\right).$$

Hence, $B' = B$ is always a solution to the Euler equation when $B' = B > \tilde{B}'$.

Let $B > B^*$. The function $\mathcal{E}(B)$ is monotone increasing. In turn, the function $\mathcal{E}'(B')$ is monotone increasing to the left of $\tilde{B}'$. Hence, for $B > B^*$ a solution must lie at $B' > \tilde{B}'$. However, for any $B' > \tilde{B}'$,

$$\mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right) = \frac{B'}{1 - (1-\beta) B'} q\left(1 - (1-\beta) B'\right)$$

which follows from the fact that $B' > \tilde{B}' > B^\star\left(1/\beta, \tilde{B}'\right)$ in which case $\Pi\left(B'; B, \tilde{B}, \tilde{B}'\right) = 1$. This implies that for any $B > B^*$ we have that:

$$\mathcal{E}\left(B; \tilde{B}\right) = \mathcal{E}'\left(B'; B, \tilde{B}, \tilde{B}'\right),$$

which implies,

$$\frac{B}{1 - (1-\beta) B} \cdot q\left(1 - (1-\beta) B\right) = \frac{B'}{1 - (1-\beta) B'} \cdot q\left(1 - (1-\beta) B'\right).$$

By monotonicity, we immediately conclude that $B = B'$, and that is the only solution.

Now consider $B \in \left(\tilde{B}', B^*\right)$. We now look for a solution $B' \in \left(B^\star(\tilde{B}'), \tilde{B}'\right)$. If $B < B^*$, again by monotonicity:

$$\mathcal{E}'\left(B; \tilde{B}, \tilde{B}, \tilde{B}'\right) < \lim_{B' \uparrow \tilde{B}'} \mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right).$$

We make use of the following calculations:

$$\begin{aligned}\mathcal{E}'\left(B';\tilde{B},\tilde{B},\tilde{B}'\right) &= \lim_{B'\downarrow\tilde{B}'}\mathcal{E}'\left(B';\tilde{B},\tilde{B},\tilde{B}'\right)\Rightarrow\\ &= \frac{B'}{1-(1-\beta)B'}\cdot\frac{q\left(1-(1-\beta)\tilde{B}'\right)}{\lim_{B'\uparrow\tilde{B}'}\Pi\left(B';\tilde{B},\tilde{B},\tilde{B}'\right)}\\ &= \frac{\tilde{B}'}{1-(1-\beta)\tilde{B}'}\cdot\frac{q\left(1-(1-\beta)\tilde{B}'\right)}{\frac{q\left(1-(1-\beta)\tilde{B}'\right)}{q\left(1-(1-\beta)\tilde{B}\right)}}\\ &= \frac{\tilde{B}'}{1-(1-\beta)\tilde{B}'}\cdot q\left(1-(1-\beta)\tilde{B}\right)>0.\end{aligned}$$

Fix $\tilde{B},\tilde{B}'$. Define $B^*>\tilde{B}$ to be the value that solves:

$$\mathcal{E}'\left(B^*;\tilde{B},\tilde{B},\tilde{B}'\right)=\lim_{B'\uparrow\tilde{B}'}\mathcal{E}'\left(B';\tilde{B},\tilde{B},\tilde{B}'\right).$$

By definition, if indeed $B^*>\tilde{B}$, then

$$\frac{B^*}{1-(1-\beta)B^*}\cdot q\left(1-(1-\beta)B^*\right)=\lim_{B'\uparrow\tilde{B}'}\frac{B'}{1-(1-\beta)B'}\cdot\frac{Q\left(B',\tilde{B}'\right)}{\Pi\left(B';B^*,\tilde{B},\tilde{B}'\right)}$$

which implies

$$B^*=\frac{\tilde{B}'}{\frac{1-(1-\beta)\tilde{B}'}{q\left(1-(1-\beta)\tilde{B}'\right)}+(1-\beta)\tilde{B}'}=\frac{\tilde{B}'}{C^w\left(\tilde{B}',\tilde{B}'\right)+C^s\left(\tilde{B}'\right)}>\tilde{B}'.$$

The last inequality verifies that $B^*>\tilde{B}'$. This result follows from

$$\lim_{B'\uparrow\tilde{B}'}\Pi\left(B';B^*,\tilde{B},\tilde{B}'\right)=\left(q\left(1-(1-\beta)B^*\right)\right)^{-1}$$

and

$$\lim_{B'\uparrow\tilde{B}'}Q\left(B',\tilde{B}'\right)=q\left(1-(1-\beta)\tilde{B}'\right).$$

The inequality holds since $q>1$. Furthermore, this satisfies $B^*>\tilde{B}$ as required since the SBL sequence is weakly increasing.[59]

[59]This is important because the relevant points to study are those for which the all chained equilibrium is possible. These are the cases where $B^*>\tilde{B}'\geq\tilde{B}$.

Next, we show that for any $B \in \left(\tilde{B}', B^*\right)$, the equation $\mathcal{E} = \mathcal{E}'$ has two roots; one above $\tilde{B}'$ and one below. The solution above $\tilde{B}'$ is trivial: $B = B' > \tilde{B}' \geq \tilde{B}$ satisfies the equation. For the solution below, we use continuity and monotonicity to show that there is a unique $B' < \tilde{B}'$ such that

$$\mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right) = \lim_{B' \downarrow \tilde{B}'} \mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right)$$

$$\frac{B'}{1-(1-\beta)B'} \cdot Q\left(B', \tilde{B}'\right) = \frac{B'}{1-(1-\beta)B'} \cdot \frac{q\left(1-(1-\beta)\tilde{B}'\right)}{\lim_{B' \uparrow \tilde{B}'} \Pi\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right)}$$

$$\frac{B'}{1-(1-\beta)B'} \cdot Q\left(B', \tilde{B}'\right) = \frac{\tilde{B}'}{1-(1-\beta)\tilde{B}'} \cdot \frac{q\left(1-(1-\beta)\tilde{B}'\right)}{\frac{q\left(1-(1-\beta)\tilde{B}'\right)}{q\left(1-(1-\beta)\tilde{B}\right)}}$$

$$\frac{B'}{1-(1-\beta)B'} \cdot Q\left(B', \tilde{B}'\right) = \frac{\tilde{B}'}{1-(1-\beta)\tilde{B}'} \cdot q\left(1-(1-\beta)\tilde{B}\right) > 0.$$

Lets call this B', $\underline{B}$ and think of the interval $\left(\tilde{B}', B^*\right)$. The interpretation of $\underline{B}$ is that it is the small root for the debt level "$\tilde{B}' + \varepsilon$". Meaning, if $\tilde{B}'$ is the small root for B^*(the end of the interval) then $\underline{B}$ is the small root for the start of the interval. Now, notice that the RHS is a constant with respect to $\tilde{B}'$. To establish existence of $\underline{B}$ further notice that the LHS tends to $\lim_{B' \uparrow \tilde{B}'} \mathcal{E}'\left(B'; B^*, \tilde{B}, \tilde{B}'\right)$ as $B' \uparrow \tilde{B}'$ (this was shown in the previous step), and this is larger than the RHS just by comparing magnitudes ($\tilde{B}' \geq \tilde{B}$ and $q\left(1-(1-\beta)B^*\right) > 1$). Then notice that as $B' \downarrow 0$ the LHS goes to zero, which is lower than the RHS. By continuity of the LHS, we can apply the intermediate value theorem for existence. Uniqueness is granted since the LHS is increasing in B'. The fraction is clearly increasing in B', the average price too because as $B' \uparrow \tilde{B}'$ the share of chained expenditure increases, and also its price does so. This statement also uses the fact that Q is the weighted harmonic mean of prices with expenditure weights. Since the LHS is continuous and increasing it maps the interval $\left(\underline{B}, \tilde{B}'\right)$ onto

$$\left(\lim_{B' \downarrow \tilde{B}'} \mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right), \lim_{B' \uparrow B^*} \mathcal{E}'\left(B'; \tilde{B}, \tilde{B}, \tilde{B}'\right)\right).$$

So (since $\mathcal{E}'$ is increasing and continuous) it covers all the image of $\left(\tilde{B}', B^*\right)$. This proves that for $B \in \left(\tilde{B}', B^*\right)$ there exist two roots that solve $\mathcal{E} = \mathcal{E}'$. One is $B_2' = B > \tilde{B}'$ and the other is a $B_1' \in \left(\underline{B}, \tilde{B}'\right)$.

To make affairs clearer, suppose that $B^* < B^h$ then for all $B \in \left(\tilde{B}', B^*\right)$ the larger

root $B' = B$ is not an equilibrium and there is a region $\left(B^*, B^h\right)$ that does not have a symmetric competitive equilibrium because the only root that solves equation (7) provides allocations and prices that are not an equilibrium. Now suppose that $B^* > B^h$, then for $B \in \left(\tilde{B}', B^h\right)$ the larger root $B' = B$ is not an equilibrium and for $B \in \left(B^h, \overline{B}\right)$. Summarizing both cases. For $B_0 < B^h\left(\tilde{B}_0\right)$, if a (symmetric competitive) equilibrium exists, then it is given by the smaller root of equation (7).

□

E.4 Proof of Corollary 1

Part i. At $B_t \leq B^\star\left(\beta^{-1}, \tilde{B}_{ss}\right)$ if $R = \beta^{-1}$ then marginal and average prices are equal to 1. Then it is evident that $B_{t+1} = B_t$ solves equation (7) and the expenditure is $1 - (1-\beta) B_t \leq \tilde{B}_{ss} - B_t$ by assumption. As a consequence, the steady state is non-disrupted.

Part ii. This case follows immediately from Proposition 3. If $B_0 > B^h$, there exists an equilibrium with $R_t = 1/\beta$ and $B_t = B_0$. Recall that if $R_t = 1/\beta$, and $B_t = B_{t+1}$ is a solution to $\mathcal{E}_t = \mathcal{E}_{t+1}$. Since for $R_t = 1/\beta$, $B_0 = B_t$ is a solution to the worker's problem, both Euler equations are satisfied, and the asset clearing condition holds, sequences with $R_t = 1/\beta$ and $B_t = B_0$ are consistent with the equilibrium definition. Since at $B_t > B^h > B^\star$, $q_t > 1$. Hence, the economy remains permanently disrupted.

E.5 Proof of Corollary 2

That $B_{t+1} < B_t$ if $B_t \in \left(\tilde{B}, B^*\left(\tilde{B}\right)\right)$ is immediate from Lemma 14 since we are choosing the smaller root and the larger root is $B_{t+1} = B_t$. For $B_t \in \left(B^\star, \tilde{B}_t\right]$, it is enough to show that $\beta R_{t+1} < 1$ as done in step 3 of the proof of Lemma 14.

F Constrained Efficiency Analysis (Section 5)

F.1 Characterization of the Primal Problem

This appendix provides the complete characterization of the Primal Planner Problem in subsection 5.3. For convenience, we reproduce the problem here and then derive the solution across different regimes of the spot borrowing limit $\tilde{B}$.

Since the maximization is static (time-separable), we can solve the problem period by period. The Primal Planner's problem at each date is:

Problem 8. (Static Primal Problem):

$$\mathcal{P}^{\theta}\left(\tilde{B}\right) = \max_{B\in\left[0,\bar{B}\right]} \mathcal{P}\left(B,\tilde{B}\right)$$

where

$$\mathcal{P}\left(B,\tilde{B}\right) = (1-\theta)\log\left((1-\beta)B\right) + \theta\log\left(\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)\mu\left(B,\tilde{B}\right) + S\left(B,\tilde{B}\right)\right)$$

subject to:

$$\mu\left(B,\tilde{B}\right) \equiv 1-(1-\beta)B - \min\left\{\max\left\{\tilde{B}-B,0\right\},1-(1-\beta)B\right\}$$

$$S\left(B,\tilde{B}\right) = \min\left\{\max\left\{\tilde{B}-B,0\right\},1-(1-\beta)B\right\}.$$

This is a problem in which the planner chooses debt for agents, while respecting the SBL and payment constraints. Then, the optimal path of debt for the planner's problem is $B_t = \mathcal{B}^{\theta}\left(\tilde{B}_t\right)$. We proceed to solve for $\mathcal{B}^{\theta}$. Welfare depends on whether and how the planner may want to distort TFP to increase egalitarian objectives, given its Pareto weight. The tradeoffs depend on the SBL, $\tilde{B}$.

We make use of two auxiliary problems, the best and worst value problems:

$$\underline{\mathcal{P}}^{\theta} \equiv \mathcal{P}^{\theta}(0) \text{ and } \bar{\mathcal{P}}^{\theta} \equiv \mathcal{P}^{\theta}\left(\bar{B}\right).$$

That is, $\underline{\mathcal{P}}^{\theta}$ is the value corresponding to the case where the entrepreneur must only make chained expenditures and $\bar{\mathcal{P}}^{\theta}$ the corresponding problem with no SBL constraints.

Recall that the unconstrained solution to the Primal Planner's problem is given by

the ratio of Pareto weights:

$$\frac{1-(1-\beta)\,B^o}{(1-\beta)\,B^o}=\frac{\theta}{1-\theta}\rightarrow B^o=\frac{1-\theta}{1-\beta}.$$

The solution is as follows:

Proposition 8. (Solution of the Primal Problem): *Let the solution to the primal problem be $\mathcal{B}^\theta$. The function $\mathcal{B}^\theta$ satisfies:*

I. Efficiency Threshold. *For $\tilde{B}\geq\frac{1-\theta\beta}{1-\beta}$, $\mathcal{B}^\theta=B^o$. Moreover, for this debt level $X=0$.*

II. Constrained Region. *For $\tilde{B}<\frac{1-\theta\beta}{1-\beta}$, the unconstrained optimum is infeasible, but productive efficiency may be preserved by distorting the consumption distribution. The solution to the Primal Planner's problem in this region depends on the threshold SBL, $\tilde{B}^i$.*

II.a Social Insurance and Productive Efficiency. *For $\tilde{B}\in\left[\tilde{B}^i,\frac{1-\theta\beta}{1-\beta}\right]$, we have that $\mathcal{B}^\theta=B^\star\left(\tilde{B}\right)$:*

$$\frac{1-(1-\beta)\,B^\star\left(\tilde{B}\right)}{(1-\beta)\,B^\star\left(\tilde{B}\right)}\leq\frac{\theta}{1-\theta}\frac{1-\beta\delta}{1-\beta}.$$

II.b Social Insurance Complements Productive Efficiency. *If $\tilde{B}^s<\tilde{B}^i$, for $\tilde{B}\in\left[\tilde{B}^s,\tilde{B}^i\right]$, we have that $\mathcal{B}^\theta$ solves:*

$$\frac{1-(1-\beta)\,B^\theta}{(1-\beta)\,B^\theta}=\frac{\theta}{1-\theta}\frac{Q\left(B^\theta,\tilde{B}\right)}{q\left(B^\theta,\tilde{B}\right)}\left(\frac{q\left(B^\theta,\tilde{B}\right)-\beta\left(1+\epsilon_\mu^{\mathcal{A}}\left(\mu\left(B^\theta,\tilde{B}\right)\right)\right)}{1-\beta}\right).$$

Moreover, for this debt level $X,S>0$.

II.c Social Insurance Conflicts Productive Efficiency. *For $\tilde{B}\in\left[0,\tilde{B}^s\right]$, we have that $\mathcal{B}^\theta$ is the unique constant solution $B^\theta>B_{ss}$ to the equation:*

$$\frac{1-(1-\beta)\,B^\theta}{(1-\beta)\,B^\theta}=\frac{\theta}{1-\theta}\left(1+\epsilon_\mu^{\mathcal{A}}\left(1-(1-\beta)\,B^\theta\right)\right).$$

Moreover, for this debt level $S=0$.

Threshold value $\tilde{B}^i$. *$\tilde{B}^i$ is the unique solution to:*

$$\frac{1-(1-\beta)\,B^\star\left(\tilde{B}^i\right)}{(1-\beta)\,B^\star\left(\tilde{B}^i\right)}=\frac{\theta}{1-\theta}\frac{1-\beta\delta}{1-\beta}.$$

Threshold value $\tilde{B}^s$. *Let $\underline{\mathcal{P}}^\theta\equiv\mathcal{P}^\theta(0)$ and $\bar{\mathcal{P}}^\theta\equiv\mathcal{P}^\theta\left(\bar{B}\right)$. The threshold $\tilde{B}^s$ solves:*

$$\bar{\mathcal{P}}^\theta=\underline{\mathcal{P}}^\theta+\int_{\tilde{B}^s}^{\frac{1-\theta\beta}{1-\beta}}\left(\mathcal{P}_{\tilde{B}}+\mathcal{P}_B|_{B=B^\star(\tilde{B})}\cdot B^\star_{\tilde{B}}\left(\tilde{B}\right)\right)d\tilde{B}$$

where for $\tilde{B}\in\left[\tilde{B}^i,\frac{1-\theta\beta}{1-\beta}\right]$ we have that $\mathcal{P}^\theta_{\tilde{B}}\left(\tilde{B}\right)$ is given in the proof.

F.2 Proof of Proposition 8 (Solution of the Primal Problem)

CaseI. Values of $\tilde{B}$ such that all consumption is spot. Ideally, the planner wants to maximize spot consumption by setting $\mu=0$. The unconstrained solution to the Primal Planner's problem sets $B=B^o$ and yields the value as $\bar{\mathcal{P}}^\theta$. With constraints, this value is achieved if and only if all spot consumption is feasible:

$$\max\left\{\tilde{B}-B^o,0\right\}\geq 1-(1-\beta)\,B^o>0.$$

Thus, we need

$$\tilde{B}\geq B^o$$

and that:

$$\tilde{B}\geq 1+\beta B^o.$$

Combining both constraints, we have:

$$\tilde{B}\geq B^o+\max\left\{1-(1-\beta)\,B^o,0\right\}.$$

We know that the optimal debt $B^o<(1-\beta)^{-1}$, i.e., it must be less than the natural borrowing limit. Hence, we get rid of the max operator:

$$\tilde{B}\geq 1+\beta B^o=1+\frac{\beta}{1-\beta}(1-\theta)=\frac{1-\theta\beta}{1-\beta}. \tag{81}$$

Thus, for these levels of the SBL that satisfy the condition above, the planner can achieve the unconstrained first best.

Case 2. Constrained allocations. Now consider the case where the constraint binds, $\tilde{B} < \frac{1-\theta\beta}{1-\beta}$. In this case, the planner cannot achieve the unconstrained solution.[60] Chained expenditures are given by:

$$\mu\left(B,\tilde{B}\right) = 1-(1-\beta)B - \min\left\{\max\left\{\tilde{B}-B,0\right\}, 1-(1-\beta)B\right\} \geq 0.$$

We have critical values.

- If the planner chooses $B^p \geq \tilde{B}$, there is no spot consumption.
- If the planner chooses $B^p < B_l\left(\tilde{B}\right) < \tilde{B}$, there is no chained consumption. The threshold $B_l\left(\tilde{B}\right)$ solves,

$$\max\left\{\tilde{B}-B_l,0\right\} = 1-(1-\beta)B_l.$$

  Since $B_l < \tilde{B}$, we obtain:

$$\tilde{B} = 1+\beta B_l \rightarrow B_l\left(\tilde{B}\right) = \max\left\{0,\beta^{-1}\left(\tilde{B}-1\right)\right\} = B^\star\left(\tilde{B}\right).$$

  $B^\star\left(\tilde{B}\right) < B^o$ since we are in the constrained region.

Consider now the planner problem that restricts choices to at least some of both goods that are consumed by the entrepreneur:

Problem 9. *The Primal Planner's problem restricted to both types of consumption is:*

$$\tilde{\mathcal{P}}\left(\tilde{B}\right) = \max_B\left\{(1-\theta)\log\left((1-\beta)B\right) + \theta\log\left(\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)\cdot\mu\left(B,\tilde{B}\right) + S\left(B,\tilde{B}\right)\right)\right\}$$

subject to:

$$B \in \left[B^\star(\tilde{B}),\tilde{B}\right]$$

$$\mu\left(B,\tilde{B}\right) \equiv 1+\beta B-\tilde{B}$$

and

$$S\left(B,\tilde{B}\right) = \tilde{B}-B.$$

We have the following Lemma.

[60] This does not necessarily require chained purchases: the planner may preserve productive efficiency by changing the consumption distribution.

Lemma 15. *For any $\tilde{B} < \frac{1-\theta\beta}{1-\beta}$, the original planner problem satisfies:*

$$\mathcal{P}^{\theta}\left(\tilde{B}\right) = \max\left\{\tilde{\mathcal{P}}\left(\tilde{B}\right), \underline{\mathcal{P}}^{\theta}\right\}.$$

Proof. Indeed, in the region $B^p \in \left[0, B^{\star}(\tilde{B})\right]$ the objective of the planner is equivalent to the objective when the SBL is most relaxed, $\mathcal{P}\left(B, \bar{B}\right)$. Thus, since $\tilde{B} < \frac{1-\theta\beta}{1-\beta}$, the planner's objective is increasing in the region $\left[0, B_l\left(\tilde{B}\right)\right]$. Thus, the planner's solution must fall in between $B^p \in \left[B_l\left(\tilde{B}\right), \bar{B}\right]$. For any $B^p \geq B_h\left(\tilde{B}\right) = \tilde{B}$, the objective function in $\mathcal{P}\left(B, \tilde{B}\right)$ is independent of $\tilde{B}$ and hence, must coincide with the value of $\underline{\mathcal{P}}^{\theta}$. Hence, we can partition $\mathcal{P}^{\theta}\left(\tilde{B}\right)$ according to the Lemma. □

To prove the main result, we solve the problems $\underline{\mathcal{P}}^{\theta}$ and $\tilde{\mathcal{P}}\left(\tilde{B}\right)$.

Auxiliary Problem $\underline{\mathcal{P}}^{\theta}$: no spot consumption. The planner's problem with the tightest SBL $\underline{\mathcal{P}}^{\theta} = \mathcal{P}^{\theta}(0)$ is given by:

$$\underline{\mathcal{P}}^{\theta} = \max_{B \in \left[0, \bar{B}\right]} \underline{\mathcal{P}}(B, 0)$$

where

$$\underline{\mathcal{P}}^{\theta} = \max_{B \in \left[0, \bar{B}\right]} \left\{(1-\theta)\log\left((1-\beta)B\right) + \theta\log\left(\mathcal{A}\left(\mu\left(B,0\right)\right)\mu\left(B,0\right)\right)\right\}$$

subject to:

$$\mu(B,0) \equiv 1 - (1-\beta)B.$$

To solve this problem, we perform some calculations. First, note that:

$$\frac{\partial\left[\mathcal{A}(\mu)\mu\right]}{\partial\mu} = \mathcal{A}(\mu)\left(1+\epsilon^{\mathcal{A}}\right)$$

where,

$$\epsilon_{\mu}^{\mathcal{A}} \equiv \frac{\partial\mathcal{A}(\mu)}{\partial\mu}\frac{\mu}{\mathcal{A}(\mu)}.$$

The derivative $\underline{\mathcal{P}}_B$ is therefore given by:

$$\underline{\mathcal{P}}_B = (1-\theta)\frac{(1-\beta)}{(1-\beta)B} + \theta\frac{\mathcal{A}(\mu)\left(1+\epsilon_{\mu}^{\mathcal{A}}\right)\mu_B(B,0)}{\mathcal{A}(\mu)\mu} = \frac{(1-\theta)}{B} - \theta\frac{\left(1+\epsilon_{\mu}^{\mathcal{A}}\right)(1-\beta)}{1-(1-\beta)B}.$$

The second equation uses that $\mu_B(B,0) \equiv -(1-\beta)$.

The first term in $\underline{\mathcal{P}}_B$, $(1-\theta)/B$, is decreasing in B. The second term,

$$\theta \frac{\left(1+\epsilon_{\mu}^{\mathcal{A}}\right)(1-\beta)}{1-(1-\beta)B}, \tag{82}$$

is increasing. We know this because the denominator is decreasing in $B \in \left[0, \bar{B}\right]$ and the elasticity of TFP is itself decreasing in μ,

$$\epsilon_{\mu\mu}^{\mathcal{A}} = \frac{\partial}{\partial \mu}\left[\frac{\mathcal{A}'(\mu)\mu}{\mathcal{A}(\mu)}\right] = \frac{\mathcal{A}''(\mu)\mu}{\mathcal{A}(\mu)} + \frac{\mathcal{A}'(\mu)}{\mathcal{A}(\mu)} - \frac{\left[\mathcal{A}'(\mu)\right]^2}{\left[\mathcal{A}(\mu)\right]^2}\mu < 0.$$

Hence, $\epsilon_{\mu\mu}^{\mathcal{A}}\mu_B(B,0) > 0$, since the product of two negative numbers is positive, thus the numerator of the second term (82) is increasing in B. Thus, $\underline{\mathcal{P}}$ is concave and therefore $\underline{\mathcal{P}}^{\theta}$ has a unique solution:

$$\frac{1-(1-\beta)B}{(1-\beta)B} = \frac{\theta}{1-\theta}\left(1+\epsilon_{\mu}^{\mathcal{A}}(\mu(B,0))\right)$$

and recall that $\mu(B,0) = 1-(1-\beta)B$. We call this solution $\underline{B}^p$: the planner debt level under the most tight SBL. We have the following Lemma:

Lemma 16. *The solution* $\underline{B}^p > B^o$.

Proof. The proof is immediate from $1+\epsilon_{\mu}^{\mathcal{A}} < 1$ and the fact that

$$\frac{1-(1-\beta)B^o}{(1-\beta)B^o} = \frac{\theta}{(1-\theta)}.$$

□

Next, we solve $\tilde{\mathcal{P}}\left(\tilde{B}\right)$.

Auxiliary Problem $\tilde{\mathcal{P}}\left(\tilde{B}\right)$: spot and chained consumption. Consider now the planner problem where at least some of both goods are consumed by the entrepreneur:

$$\tilde{\mathcal{P}}\left(\tilde{B}\right) = \max_{B \in \left[B^{\star}\left(\tilde{B}\right), \tilde{B}\right]} \tilde{\mathcal{P}}\left(B, \tilde{B}\right)$$

where

$$\tilde{\mathcal{P}}\left(B,\tilde{B}\right) \equiv (1-\theta)\log\left((1-\beta)B\right) + \theta\log\left(\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)\cdot\mu\left(B,\tilde{B}\right) + S\left(B,\tilde{B}\right)\right),$$

subject to:

$$\mu\left(B,\tilde{B}\right)\equiv 1+\beta B-\tilde{B}$$

and

$$S\left(B,\tilde{B}\right)=\tilde{B}-B.$$

The derivative of the objective in $\tilde{\mathcal{P}}\left(\tilde{B}\right)$ is:

$$\tilde{\mathcal{P}}_B\left(B,\tilde{B}\right)=(1-\theta)\frac{1}{B}+\theta\frac{\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)\left(1+\epsilon_\mu^{\mathcal{A}}\left(\mu\left(B,\tilde{B}\right)\right)\right)\beta-1}{\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)\mu\left(B,\tilde{B}\right)+\tilde{B}-B}.$$

Recall that,

$$C\left(B,\tilde{B}\right)=\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)\mu\left(B,\tilde{B}\right)+\tilde{B}-B$$

and

$$Q\cdot C\left(B,\tilde{B}\right)=E\left(B,\tilde{B}\right)=1-(1-\beta)B.$$

Hence, using the definition of Q and q we rewrite:

$$\tilde{\mathcal{P}}_B\left(B,\tilde{B}\right)=(1-\theta)\frac{1}{B}-\theta Q\frac{1-\beta\left(1+\epsilon_\mu^{\mathcal{A}}\left(\mu\left(B,\tilde{B}\right)\right)\right)\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)}{1-(1-\beta)B}.$$

We can multiply both sides by the ratio of $1-(1-\beta)B$ and divide by $(1-\beta)(1-\theta)$ and obtain:

$$\tilde{\mathcal{P}}_B\left(B,\tilde{B}\right)\frac{1-(1-\beta)B}{(1-\theta)(1-\beta)}=\frac{1-(1-\beta)B}{(1-\beta)B}-\frac{\theta}{1-\theta}Q\left(\frac{1-\beta\left(1+\epsilon_\mu^{\mathcal{A}}\right)\mathcal{A}(\mu)}{1-\beta}\right).$$

This function must have the same sign as $\tilde{\mathcal{P}}_B\left(B,\tilde{B}\right)$, since it was obtained by multiplication of positive numbers. The first term is decreasing in B. In turn, $Q_\mu\mu_B$ is increasing in B. Hence, as long as

$$\mathcal{A}(\mu)\left(1+\epsilon_\mu^{\mathcal{A}}\right)=\mathcal{A}(\mu)+\mathcal{A}'(\mu)\mu$$

is decreasing in B, the second term is increasing. The second term is indeed decreasing in μ since its derivative is:

$$2\mathcal{A}'(\mu)+\mathcal{A}''(\mu)\mu<0,$$

where the sign follows immediately from the concavity and monotone decreasing prop-

erties of $\mathcal{A}$ (equivalently, since $2\mathcal{A}' + \mu\mathcal{A}'' = \mathcal{Y}''$, from the concavity of $\mathcal{Y}$). Hence, the objective function $\tilde{\mathcal{P}}\left(B, \tilde{B}\right)$ is concave in B. Furthermore, since

$$Q\left(\frac{1-\beta\left(1+\epsilon_{\mu}^{\mathcal{A}}\right)\mathcal{A}(\mu)}{1-\beta}\right) > 1,$$

and we have an interior maximum in the region $B \in \left[B^{\star}\left(\tilde{B}\right), \tilde{B}\right]$, B must solve:

$$\frac{1-(1-\beta)B}{(1-\beta)B} = \frac{\theta}{1-\theta}Q\left(\frac{1-\beta\left(1+\epsilon_{\mu}^{\mathcal{A}}\right)\mathcal{A}(\mu)}{1-\beta}\right),$$

and must be such that $B < B^{o}$.

Next, we establish properties regarding the limits of this function at $B^{\star}\left(\tilde{B}\right)$ and $\tilde{B}$. We start with $B^{\star}\left(\tilde{B}\right)$:

$$\lim_{B\downarrow B^{\star}\left(\tilde{B}\right)} \frac{1-(1-\beta)B}{(1-\beta)B} - \frac{\theta}{1-\theta}Q\left(\frac{1-\beta\left(1+\epsilon_{\mu}^{\mathcal{A}}\right)\mathcal{A}(\mu)}{1-\beta}\right) =$$

$$\frac{1-(1-\beta)B^{\star}\left(\tilde{B}\right)}{(1-\beta)B^{\star}\left(\tilde{B}\right)} - \frac{\theta}{1-\theta}\left(\frac{1}{1-\beta} - \frac{\beta\left(1+\lim_{\mu\downarrow 0}\epsilon_{\mu}^{\mathcal{A}}\right)\lim_{\mu\downarrow 0}\mathcal{A}(\mu)}{1-\beta}\right),$$

where we used $\lim_{\mu\downarrow 0} Q = 1$. Appendix A, shows that $\lim_{\mu\downarrow 0}\left(1+\epsilon_{\mu}^{\mathcal{A}}\right) = 1$ and that $\lim_{\mu\downarrow 0}\mathcal{A}(\mu) = \delta$. Thus, the limit of the objective function at the left boundary has the sign of:

$$\frac{1-(1-\beta)B^{\star}\left(\tilde{B}\right)}{(1-\beta)B^{\star}\left(\tilde{B}\right)} - \frac{\theta}{1-\theta}\frac{1-\beta\delta}{1-\beta} = \frac{1-(1-\beta)B^{\star}\left(\tilde{B}\right)}{(1-\beta)B^{\star}\left(\tilde{B}\right)} - \frac{1-(1-\beta)B^{o}}{(1-\beta)B^{o}}\frac{1-\beta\delta}{1-\beta}.$$

The function $B^{\star}\left(\tilde{B}\right) < B^{o}$ but $\frac{1-\beta\delta}{1-\beta} > 1$, hence the sign is ambiguous. The solution is at this corner if the function is negative.

Next, we consider the limit of the derivative of the objective at $\tilde{B}$. Evaluating the limits is immediate. Hence, if

$$\frac{1-(1-\beta)\tilde{B}}{(1-\beta)\tilde{B}} \geq \frac{\theta}{1-\theta}\left(\frac{1-\beta\left(1+\epsilon_{\mu}^{\mathcal{A}}(\mu)\right)\mathcal{A}(\mu)}{1-\beta}\right)\Bigg|_{\mu=\left(1-(1-\beta)\tilde{B}\right)}$$

the solution is $B = \tilde{B}$. Moreover, we know that since

$$\left(\frac{1 - \beta \left(1 + \epsilon_{\mu}^{\mathcal{A}}\right) \mathcal{A}(\mu)}{1 - \beta} \right) > 1,$$

the corner solution $\tilde{B}$ is chosen only if

$$\frac{1 - (1 - \beta) \tilde{B}}{(1 - \beta) \tilde{B}} > \frac{1 - (1 - \beta) B^o}{(1 - \beta) B^o}.$$

Collecting the results up to this point, we have the following Lemma.

Lemma 17. *The solution B_{ℓ}^{p} to $\tilde{\mathcal{P}}\left(\tilde{B}\right)$ is as follows.*

I. $B_{\ell}^{p} = \tilde{B}$ *if*

$$\frac{1 - (1 - \beta) \tilde{B}}{(1 - \beta) \tilde{B}} \geq \frac{\theta}{1 - \theta} \left(\frac{1 - \beta \left(1 + \epsilon_{\mu}^{\mathcal{A}}(\mu)\right) \mathcal{A}(\mu)}{1 - \beta} \right) \Bigg|_{\mu = \left(1 - (1 - \beta)\tilde{B}\right)},$$

II. $B_{\ell}^{p} = B^{\star}\left(\tilde{B}\right)$ *if*

$$\frac{1 - (1 - \beta) B^{\star}\left(\tilde{B}\right)}{(1 - \beta) B^{\star}\left(\tilde{B}\right)} \leq \frac{\theta}{1 - \theta} \frac{1 - \beta\delta}{1 - \beta}.$$

III. *Otherwise, B_{ℓ}^{p} solves:*

$$\frac{1 - (1 - \beta) B_{\ell}^{p}}{(1 - \beta) B_{\ell}^{p}} = \frac{\theta}{1 - \theta} Q\left(B_{\ell}^{p}, \tilde{B}\right) \left(\frac{1 - \beta \left(1 + \epsilon_{\mu}^{\mathcal{A}}\left(\mu\left(B_{\ell}^{p}, \tilde{B}\right)\right)\right) \mathcal{A}\left(\mu\left(B_{\ell}^{p}, \tilde{B}\right)\right)}{1 - \beta} \right) \Bigg|_{\mu = \left(1 + \beta B_{\ell}^{p} - \tilde{B}\right)}.$$

The following Lemma characterizes threshold values for $\tilde{B}$ corresponding to the Lemma above:

Lemma 18. *The condition*

$$\frac{1 - (1 - \beta) B^{\star}\left(\tilde{B}\right)}{(1 - \beta) B^{\star}\left(\tilde{B}\right)} \leq \frac{\theta}{1 - \theta} \frac{1 - \beta\delta}{1 - \beta}$$

holds for all $\tilde{B} \geq \tilde{B}^i$ *such that:*

$$\frac{1-(1-\beta)\,B^{\star}\left(\tilde{B}^i\right)}{(1-\beta)\,B^{\star}\left(\tilde{B}^i\right)} = \frac{\theta}{1-\theta}\frac{1-\beta\delta}{1-\beta}.$$

The reverse inequality holds for $\tilde{B} \leq \tilde{B}^a$ *such that:*

$$\frac{1-(1-\beta)\,B^{\star}\left(\tilde{B}^a\right)}{(1-\beta)\,B^{\star}\left(\tilde{B}^a\right)} = \frac{\theta}{1-\theta}\left(\frac{1-\beta\left(1+\epsilon_{\mu}^{\mathcal{A}}(\mu)\right)\mathcal{A}(\mu)}{1-\beta}\right)\Bigg|_{\mu=\left(1-(1-\beta)\tilde{B}^a\right)}.$$

Proof. The proof follows immediately from the fact that:

$$\frac{1-(1-\beta)\,B^{\star}\left(\tilde{B}\right)}{(1-\beta)\,B^{\star}\left(\tilde{B}\right)},$$

is decreasing in $\tilde{B}$ and the function:

$$\frac{\theta}{1-\theta}Q\left(B,\tilde{B}\right)\left(\frac{1-\beta\left(1+\epsilon_{\mu}^{\mathcal{A}}\left(\mu\left(B,\tilde{B}\right)\right)\right)\mathcal{A}\left(\mu\left(B,\tilde{B}\right)\right)}{1-\beta}\right)\Bigg|_{\mu=(1-(1-\beta)B)}$$

increasing in B. □

Overall Solution. We showed above that:

$$\mathcal{P}^{\theta}\left(\tilde{B}\right) = \max\left\{\tilde{\mathcal{P}}\left(\tilde{B}\right), \underline{\mathcal{P}}^{\theta}\right\}.$$

The following Lemma shows that the solution B^p to the Planner's problem is never at $B^p = \tilde{B}$.

Lemma 19. *The planner never chooses* $B^p = \tilde{B}$.

Proof. To prove this Lemma, observe that the left limit as $B \uparrow \tilde{B}$ satisfies

$$\frac{q-1-\epsilon^{\mathcal{A}}}{1-\beta} + 1 + \epsilon^{\mathcal{A}} > 1 + \epsilon^{\mathcal{A}},$$

where the inequality follows from $q > 1$ and $\epsilon^{\mathcal{A}} < 0$. As a consequence,

$$\frac{1-(1-\beta)\,\tilde{B}}{(1-\beta)\,\tilde{B}} \geq \frac{\theta}{1-\theta}\lim_{B\uparrow\tilde{B}}\Lambda\left(B,\tilde{B}\right)$$

implies

$$\frac{1-(1-\beta)\,\tilde{B}_t}{(1-\beta)\,\tilde{B}_t} > \frac{\theta}{1-\theta}\lim_{B\downarrow\tilde{B}_t}\Lambda\left(B,\tilde{B}\right).$$

Hence, although the derivative of the objective is discontinuous at $\tilde{B}$, we know that the derivative is weakly increasing from the left and increasing from the right. This implies that $B = \tilde{B}$ is never an optimal choice. □

Next, observe $\tilde{\mathcal{P}}\left(\tilde{B}\right)$ has a compact-valued and continuous constraint correspondence with a continuous objective. Hence, it satisfies the conditions for the Theorem of the Maximum. In addition, it is immediate to verify that:

$$\lim_{\tilde{B}\uparrow\frac{1-\theta\beta}{1-\beta}}\tilde{\mathcal{P}}\left(\tilde{B}\right) = \bar{\mathcal{P}}^{\theta} \text{ and } \lim_{\tilde{B}\uparrow\frac{1-\theta\beta}{1-\beta}} B^p(\tilde{B}) = B^o.$$

We employ the Envelope Theorem on $\tilde{\mathcal{P}}\left(\tilde{B}\right)$. In the region where the solution to $\tilde{\mathcal{P}}\left(\tilde{B}\right)$ is not at a corner solution, the Envelope Theorem yields:

$$\tilde{\mathcal{P}}_{\tilde{B}} = \frac{\theta}{1-(1-\beta)\,B^p(\tilde{B})}Q\left(B^p(\tilde{B}),\tilde{B}\right)\left(1-\left(1+\epsilon^{\mathcal{A}}_{\mu}\left(\mu\left(B^p(\tilde{B}),\tilde{B}\right)\right)\right)\mathcal{A}\left(\mu\left(B^p(\tilde{B}),\tilde{B}\right)\right)\right) > 0,$$

since $\tilde{B}$ appears directly through μ and S in the objective.

The function is strictly increasing in $\tilde{B}$ since $\epsilon^{\mathcal{A}}_{\mu} < 0$ and $\mathcal{A} < 1$. In the region where the function is at the lower corner of the constraint:

$$B = B^{\star}\left(\tilde{B}\right),$$

the value of this term is:

$$\tilde{\mathcal{P}}_{\tilde{B}} = \frac{\theta}{1-(1-\beta)\,B^{\star}\left(\tilde{B}\right)}(1-\delta) > 0$$

and

$$\tilde{\mathcal{P}}_{B}|_{B=B^{\star}(\tilde{B})} = \frac{1-\theta}{(1-\beta)\,B^{\star}\left(\tilde{B}\right)} - \frac{\theta\,(1-\beta\delta)}{1-(1-\beta)\,B^{\star}\left(\tilde{B}\right)}.$$

Thus, we have that the marginal objective value is:

$$\tilde{\mathcal{P}}_{\tilde{B}} + \tilde{\mathcal{P}}_{B}\left(B^{p}(\tilde{B})\right) B^{\star}_{\tilde{B}}\left(\tilde{B}\right) = \frac{1-\theta}{(1-\beta)\, B^{\star}\left(\tilde{B}\right)}\beta^{-1} + \frac{\theta\delta}{1-(1-\beta)\, B^{\star}\left(\tilde{B}\right)} > 0.$$

Hence, the envelope condition guarantees that $\tilde{\mathcal{P}}$ is increasing in $\tilde{B}$.

There exists a threshold $\tilde{B}^{s}$ such that:

$$\tilde{\mathcal{P}}\left(\tilde{B}^{s}\right) = \underline{\mathcal{P}}^{\theta}.$$

In particular, it solves:

$$\bar{\mathcal{P}}^{\theta} = \underline{\mathcal{P}}^{\theta} + \int_{\tilde{B}^{s}}^{\frac{1-\theta\beta}{1-\beta}} \left(\tilde{\mathcal{P}}_{\tilde{B}} + \tilde{\mathcal{P}}_{B}\left(B^{p}(\tilde{B})\right) B^{\star}_{\tilde{B}}\left(\tilde{B}\right)\right) d\tilde{B}.$$

Notice that $\underline{\mathcal{P}}^{\theta}$ is finite and lower than the value of the problem without constraints. $\tilde{\mathcal{P}}(0)$ tends to $-\infty$ and $\tilde{\mathcal{P}}\left(\overline{B}\right)$ tends to the value of problem without constraints. This implies that $\tilde{\mathcal{P}}(0) < \underline{\mathcal{P}}^{\theta} < \tilde{\mathcal{P}}\left(\overline{B}\right)$ and from continuity of the value function (Theorem of the Maximum) and the Intermediate Value Theorem, the existence of $\tilde{B}^{s}$ is guaranteed. The solution to $\tilde{B}^{s}$ follows from the fundamental theorem of calculus.

Recall that in the previous Lemma we showed that $\tilde{B}$ is never a solution to $\tilde{\mathcal{P}}$ and the monotonicity implies the existence of the threshold $\tilde{B}^{i}$. These thresholds segment the intervals as given by the proposition. This concludes the proof of Proposition 8.

F.3 Discussion of Planner's Solution.

The solution to the Primal Problem reveals a novel insight. Namely, during transitions away from payment-chain crises, debt may be inefficiently high or inefficiently low. The intuition behind this is that, for a given debt level, the planner can increase output by reducing or increasing debt. By reducing debt, the planner distributes wealth toward the worker and frees credit lines, inducing greater spot expenditures. Alternatively, by increasing debt, the planner distributes wealth toward the saver, also stimulating spot expenditures. Both distributive policies distort equality, but increase efficiency by speeding up production.

The analytic expression for B^{p} reveals how the planner allocates expenditures, balancing productive efficiency against equality differently, depending on whether $\tilde{B}$ falls in several intervals. There are four intervals of values of $\tilde{B}$ where the planner's solution is qualitatively different. These actual critical values depend on the threshold points

$\frac{1-\theta\beta}{1-\beta}$ and $\left\{\tilde{B}^s, \tilde{B}^i\right\}$ whose formulas are found in the proof.

The first interval of values is where $\tilde{B}$ is above a threshold level, above which the SBL constraint is not binding. In particular, if $\tilde{B} \geq \frac{1-\theta\beta}{1-\beta}$, the planner sets $B = B_{ss}$, and production is efficient. Furthermore, the ratio of marginal utilities equals the ratio of Pareto weights—the efficient equality condition (8) is satisfied.

If $\tilde{B}$ is below the efficiency threshold, $\frac{1-\theta\beta}{1-\beta}$, the planner distorts either equality or productive efficiency. The novelty is that the planner has two ways to increase productive efficiency: the planner can increase efficiency on the margin, either by distributing wealth toward the worker if $B < \tilde{B}$ or by distributing wealth toward the saver if $B > \tilde{B}$. To see this, observe that in the region where $B < \tilde{B}$, any further reduction in debt translates, on the margin, into more spot expenditures by the worker and, thus, increases efficiency. In the region where $B > \tilde{B}$, any increase in debt translates, on the margin, into more spot expenditures by the saver, also increasing efficiency. Because of this ambivalent nature, the planner's objective function $\mathcal{P}\left(B, \tilde{B}\right)$ is not concave in B, leading to bang-bang solutions as $\tilde{B}$ varies.

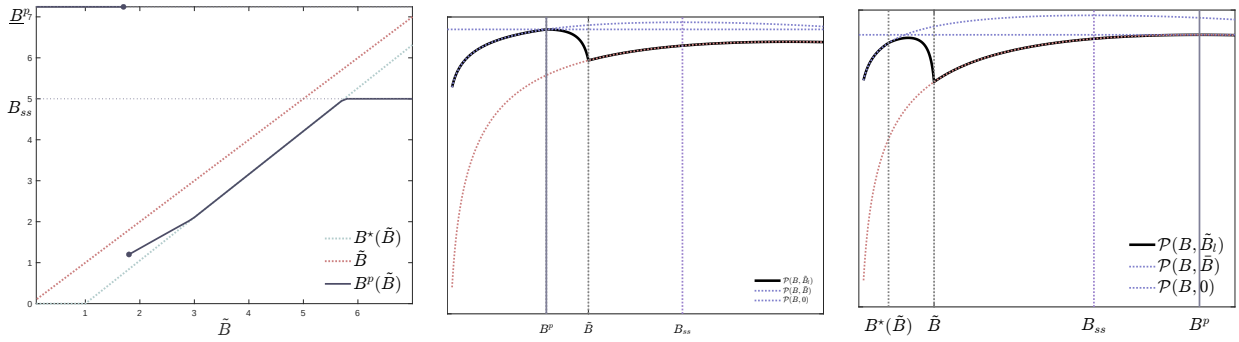
Figure F.1 describes the economics of the Primal Problem, in the regions away from efficiency. The left panel plots B^p for different values of $\tilde{B}$. In the interval of values where $\tilde{B}$ falls in a second interval, i.e., when $\tilde{B} \in [\tilde{B}^i, \frac{1-\theta\beta}{1-\beta})$, the planner solution achieves productive efficiency, $\mathcal{Y} = 1$, but the ratio of marginal utilities no longer equals the ratio of Pareto weights, as in (8). To understand why the planner does not distort production, observe that because $B^\star\left(\tilde{B}\right) < B_{ss}$, productive efficiency can only be achieved by setting $B^p = B^\star\left(\tilde{B}\right) < B_{ss}$. Setting debt at $B^\star\left(\tilde{B}\right)$, which is less than its steady-state value, implies that saver expenditures are less than in steady-state. Thus, in this interval, the planner sacrifices equality by redistributing wealth toward the worker to guarantee productive efficiency. The planner's solution is at a corner because the derivative of his objective is discontinuous at $B^\star\left(\tilde{B}\right)$. This discontinuity follows from the payments-chain network structure: the property that $\mathcal{A}(0) < 1$, which captures that even when the chained expenditure ratio is zero, an individual chained order features a production delay.

When $\tilde{B}$ falls in the third interval, $\tilde{B} \in \left[\tilde{B}^s, \tilde{B}^i\right)$, the Planner begins to sacrifice productive efficiency. To maintain productive efficiency, for those levels of the SBL, the planner would have to redistribute even more wealth to the worker at the expense of saver expenditures. In that region, sacrificing equality alone is not worth it. Hence, the planner sets debt above its efficient level, thus $B^\star\left(\tilde{B}\right) < B^p < \tilde{B}$. Therefore, in this interval, the planner still redistributes wealth (relative to steady state) toward the

agent facing the financial constraint, but the planner does partially sacrifice productive efficiency. The planner balances productive efficiency and equality.

When $\tilde{B}$ falls below an extreme value, in the fourth interval, $\tilde{B} < \tilde{B}^s$, the nature of the planner's solution changes dramatically. Increasing productive efficiency by redistributing wealth toward the worker requires the planner to set debt below $\tilde{B}$. However, when $\tilde{B} < \tilde{B}^s$, setting debt so low would require an extremely high sacrifice of saver consumption. Below that threshold, the planner prefers to redistribute wealth away from the worker, the constrained agent. Since the planner sets $B > \tilde{B}$, the planner induces only chained expenditures by the worker. Once the worker only makes chained expenditures, the SBL becomes irrelevant, so the planner chooses a constant debt level in this region. This debt level is higher than the unconstrained ideal debt level B_{ss}. Debt is higher than in steady state because the planner understands that a marginal increase in debt increases the wealth and, therefore, spot expenditures of savers. This increases productive efficiency.

The ambivalent nature of the planner's problem is central to understanding payment-chain crises. In typical models with pecuniary externalities, a planner wants to rebalance wealth toward the financially constrained agents to increase productive efficiency. Thus, typically, equality and productive efficiency reinforce each other. Here, for extremely low values of $\tilde{B}$, the planner switches to a policy mix where equality and efficiency are in conflict. The middle and right panels of Figure F.1 illustrate the change in the planner's strategy: In the middle panel, the planner prefers a value of debt where equality and efficiency complement. The right panel shows how he switches strategy as the SBL is tighter.



(a) Planner debt choice (b) $\mathcal{P}\left(\cdot,\tilde{B}\right)$ (moderate $\tilde{B}$) (c) $\mathcal{P}\left(\cdot,\tilde{B}\right)$ (tight $\tilde{B}$)

Figure F.1: Primal Planner Solution

Note: This figure reports a numerical example of the planner's solution. Figures are calculated using value function iteration: $\beta = 0.95$, $\delta = 0.9$, and $\theta = 0.75$. Panel (a) plots B^p as a function of values of $\tilde{B}$ in the range $[0.1, \theta + 0.1] * B^{star}$. Panels (b) and (c) plot the objective of the planner for different values of B in the range $[0, 1.8]$. Panel (b) fixes $\tilde{B}$ at 0.15 and Panel (c) at 0.08.

F.4 Decentralization with Taxes

To characterize the constrained efficient allocation, we demonstrate that the Ramsey problem is equivalent to a problem where a planner chooses a path of debt respecting the payments technology and the desired expenditure rules of both agents.

For ease of exposition, we reproduce both problems here. The Ramsey problem is:

Problem 10. (Ramsey Problem): *Given* B_0 *and* $\left\{\tilde{B}_t\right\}$*:*

$$\mathcal{V}^{Ramsey} = \max_{\left\{\tau^k_{t+1},\tau^c_t\right\}_{t\geq 0}} \sum_{t\geq 0} \beta^t \left[(1-\theta)\log\left(C_t\right) + \theta \log\left(S_t + X_t\right)\right],$$

subject to the creditor budget constraint and optimality:

$$\left(1+\tau^k_{t+1}\right) R^{-1}_{t+1} B_{t+1} + \left(1+\tau^c_t\right) C_t = B_t,\ \forall t \geq 0,$$

the entrepreneur budget constraint:

$$B_t + \left(1+\tau^c_t\right) E_t = R^{-1}_{t+1} B_{t+1} + Y,\ \forall t \geq 0,$$

the government's budget constraint:

$$\tau^k_{t+1} R^{-1}_{t+1} B_{t+1} + \tau^c_t \left(C_t + E_t\right) = 0,\ \forall t \geq 0,$$

and the three conditions of a competitive equilibrium: (i) Optimization, (ii) Asset and Goods market clearing, and (iii) the price consistency.

In turn, the Primal Problem is:

Problem 11. (Primal Planner Problem): *Taking* $\left\{\tilde{B}_t\right\}$ *as given:*

$$\mathcal{V}^{Primal} = \max_{\{B_t\}_{t\geq 0}} \sum_{t\geq 0} \beta^t \mathcal{P}\left(B_t, \tilde{B}_t\right)$$

where

$$\mathcal{P}\left(B,\tilde{B}\right) \equiv (1-\theta)\log\left((1-\beta)B\right) + \theta\log\left(S\left(B,\tilde{B}\right) + X\left(B,\tilde{B}\right)\right), \tag{83}$$

and taking as given $S\left(B,\tilde{B}\right)$ *and* $X\left(B,\tilde{B}\right)$ *as defined by Proposition 4.*

The Primal Problem differs from the competitive equilibrium in that the planner directly chooses the sequence of debt. The allocation and final goods is the same as the competitive equilibrium.

We will prove the following Proposition:

Proposition 9. (Primal and Ramsey Equivalence): *Let $\left\{B^p_{t+1}\right\}_{t\geq 0}$ be a solution to the Primal Planner Problem. The solution to the Ramsey problem achieves the same value as the Primal Planner Problem, $\mathcal{V}^{Ramsey} = \mathcal{V}^{Primal}$. The solution to the Primal Planner Problem can be implemented setting $(1+\tau^c_0) = B_0/B^p_0$ and a sequence of debt taxes given in the proof below.*

The proposition asserts that the solution to the Ramsey Problem can be obtained from the solution of a Primal Problem where the planner directly chooses the sequence of debt.[61]

If a solution to the primal can be implemented with taxes, it solves the Ramsey Problem. The proposition shows that this is the case. The implementation of the Primal Problem is possible because the Ramsey planner can set capital taxes to induce a desired path of debt $\{B_t\}$. Controlling B_t is key to the implementation.[62]

Before we proceed to the proof, we establish the following useful intermediate result:

Lemma 20 (Creditor's Solution with Taxes)**.** *Let $\{\tau^k_{t+1}, \tau^c_t, \tau^\ell_{t+1}\}_{t\geq 0}$ be a sequence of taxes in the Ramsey Problem. The solution to the creditor's problem is:*

$$C_t = (1-\beta)B^o_t, \quad (1+\tau^c_t)C_t = (1-\beta)B^o_t$$

where $B^o_t = B_t/(1+\tau^c_t)$ and $B^o_{t+1} = \hat{R}_{t+1}\beta B^o_t$ with:

$$\hat{R}_{t+1} \equiv \frac{R_{t+1}}{(1+\tau^k_{t+1})}\frac{1}{(1+\tau^c_{t+1})/(1+\tau^c_t)}.$$

Proof. The creditor's problem with taxes is:

$$V_t = \sum_{s=0}^{\infty} \beta^s \log(C_{t+s})$$

[61] This relation follows because the constraint set in the Primal Problem includes the constraints of the Ramsey Problem. This is immediate since market clearing in the asset market and the budget balance imply, by Walras's law, that the resource constraint holds. In the implementation, capital taxes and a period-zero expenditure tax are required by the Ramsey planner to distort the evolution of debt and replicate the solution to the Primal Problem.

[62] Time separability also implies that expenditure taxes are redundant after $t = 0$. All that the Ramsey planner needs is a sequence of capital taxes to control the path of debt and the labor tax to redistribute the capital tax receipts toward the worker. The planner only needs expenditure taxes at $t = 0$ because debt is predetermined at $t = 0$.

subject to $(1+\tau^k_{t+1})R^{-1}_{t+1}B_{t+1} + (1+\tau^c_t)C_t = B_t$. Dividing both sides by $(1+\tau^c_t)$ and defining $B^o_t \equiv B_t/(1+\tau^c_t)$, we can re-write the budget constraint as:

$$\hat{R}^{-1}_{t+1}B^o_{t+1} + C_t = B^o_t \tag{84}$$

where $\hat{R}_{t+1} \equiv \frac{R_{t+1}}{(1+\tau^k_{t+1})} \cdot \frac{1}{(1+\tau^c_{t+1})/(1+\tau^c_t)}$.

With log utility and a budget constraint like (84), the creditor's optimal expenditure rule is:

$$C_t = (1-\beta)B^o_t$$

which yields $B^o_{t+1} = \hat{R}_{t+1}\beta B^o_t$. Since $B_t = (1+\tau^c_t)B^o_t$, we have $C_t = (1+\tau^c_t)(1-\beta)B^o_t$ as claimed. □

F.5 Proof Ramsey-Primal Equivalence - Proposition 9

Proof. To characterize the constrained efficient allocation, the strategy is to show that any solution to the Ramsey Problem satisfies the constraints of the Primal Problem (Step 1). Then, we show that any solution to the Primal Problem can be induced by a proper tax sequence $\{\tau^k_{t+1}, \tau^c_t, \tau^\ell_{t+1}\}_{t\geq 0}$ (Step 2).

Step 1: From Ramsey to Primal - Deriving the Resource Constraint. Adding the two agents' budget constraints, we arrive at:

$$(1+\tau^c_t)(S_t + q_t X_t + C_t) + \tau^k_{t+1}\frac{B_{t+1}}{R_{t+1}} = 1 - \tau^\ell_{t+1}.$$

Subtracting the government budget constraint, we arrive at:

$$\begin{aligned}&(1+\tau^c_t)(S_t + q_t X_t + C_t) + \tau^k_{t+1}\frac{B_{t+1}}{R_{t+1}} \\ &\quad - \left(\tau^k_{t+1}\frac{B_{t+1}}{R_{t+1}} + \tau^c_t(C_t + S_t + q_t X_t) + \tau^\ell_{t+1}\right) = 1 - \tau^\ell_{t+1}\end{aligned}$$

which simplifies to:

$$S_t + q_t X_t + C_t = 1.$$

From Lemma 20, $C_t = (1-\beta)B^o_t$, given $B^o_t = B_t/(1+\tau^c_t)$. Therefore:

$$S_t + q_t X_t = 1 - (1-\beta)B^o_t. \tag{85}$$

Thus, setting $E_t = S_t + q_t X_t$ we have that the entrepreneurs' expenditures are:

$$E_t = 1 - (1-\beta)B_t^o. \tag{86}$$

The optimality of spot and chained expenditures is still given by the conditions in Proposition 4.

Using $q_t = \mathcal{A}(\mu_t)^{-1}$ and $\mu_t = q_t X_t = \mathcal{A}(\mu_t)^{-1} X_t$. These are exactly the constraints of the Primal Problem with B_t^o as the state variable instead of B_t^p.

Step 2: Tax Implementation of the Ramsey problem. Let $\{B_t^p\}_{t\geq 0}$ solve the Primal Problem. We now show how to construct taxes that implement this as a Ramsey equilibrium. Set the initial consumption tax:

$$(1+\tau_0^c) = B_0/B_0^p.$$

This ensures $B_0^o = B_0/(1+\tau_0^c) = B_0^p$.

For the Primal solution to be implementable, we need both agents' Euler equations to hold. The creditor's Euler equation:

$$\frac{(1-\beta)B_{t+1}^o/(1-\beta)B_t^o}{\beta} \cdot \frac{(1+\tau_{t+1}^c)}{(1+\tau_t^c)} \cdot (1+\tau_{t+1}^k) = R_{t+1} \tag{87}$$

The entrepreneur's Euler equation (using the notation from Proposition 2):

$$\frac{(S_{t+1}+X_{t+1})/(S_t+X_t)}{\beta} \cdot \frac{(1+\tau_{t+1}^c)}{(1+\tau_t^c)} \cdot \frac{\tilde{q}_{t+1}^B(B_{t+1}^o)}{\tilde{q}_t^E(B_{t+1}^o)} = R_{t+1}. \tag{88}$$

where the marginal price ratio captures the payment chain friction. Eliminating $[(1+\tau_{t+1}^c)/(1+\tau_t^c)]\beta R_{t+1}$ from both equations:

$$(1+\tau_{t+1}^k) = \frac{B_t^o}{B_{t+1}^o} \frac{S_{t+1}+X_{t+1}}{S_t+X_t} \frac{\tilde{q}_{t+1}^B(B_{t+1}^o)}{\tilde{q}_t^E(B_{t+1}^o)}.$$

Since $S_t + X_t = 1-(1-\beta)B_t^o$ and using average prices $Q_t = (S_t + q_t X_t)/(S_t + X_t)$:

$$\tau^k(B_t^o, B_{t+1}^o, \tilde{B}_t, \tilde{B}_{t+1}) = \frac{B_t^o}{B_{t+1}^o} \cdot \frac{1-(1-\beta)B_{t+1}^o}{1-(1-\beta)B_t^o} \cdot \frac{Q(B_t^o, \tilde{B}_t)}{Q(B_{t+1}^o, \tilde{B}_{t+1})} \cdot \frac{\tilde{q}_{t+1}^B(B_{t+1}^o)}{\tilde{q}_t^E(B_{t+1}^o)} - 1 \tag{89}$$

The equilibrium interest rate follows from (87):

$$R_{t+1} = \frac{(1+\tau_{t+1}^c)}{(1+\tau_t^c)} \cdot \frac{B_{t+1}^o}{\beta B_t^o} \cdot (1+\tau^k(B_t^o, B_{t+1}^o, \tilde{B}_t, \tilde{B}_{t+1})).$$

Step 3: Indeterminacy of Consumption Taxes Note that consumption taxes $\{\tau_t^c\}_{t\geq 1}$ are indeterminate. To see this, substitute the creditor's Euler equation into their budget constraint:

$$\frac{(1-\beta)B_t^o}{(1-\beta)B_{t+1}^o}\beta\left[\frac{(1+\tau_t^c)}{(1+\tau_{t+1}^c)}\right]B_{t+1}+(1+\tau_t^c)(1-\beta)B_t^o=B_t.$$

Since $B_t = (1+\tau_t^c)B_t^o$, this is satisfied for any sequence of consumption taxes. The government budget constraint then determines:

$$\tau_{t+1}^{\ell}=-\tau_{t+1}^{k}\frac{B_{t+1}}{R_{t+1}}-\tau_t^c\left((1-\beta)B_t^o+\frac{1-(1-\beta)B_t^o}{Q(B_t^o,\tilde{B}_t)}\right).$$

Setting $\tau_t^c = 0$ for all $t \geq 1$ yields the simple implementation:

$$\tau_{t+1}^{\ell}=-\frac{\tau^k(B_t^o,B_{t+1}^o,\tilde{B}_t,\tilde{B}_{t+1})}{1+\tau^k(B_t^o,B_{t+1}^o,\tilde{B}_t,\tilde{B}_{t+1})}\cdot\beta B_t^o.$$

This completes the proof that $\mathcal{V}^{Ramsey} = \mathcal{V}^{Primal}$. □

F.6 Derivation of the Bocola Effect - Proof of Proposition 5

Let's define the indirect social utility function of the Planner Problem with expenditures $\mathcal{P}^g$:

$$\mathcal{P}^g\left(B,\tilde{B},G^s,G^x\right)=(1-\theta)\log\left(C^s\left(B\right)\right)+\theta\log\left(S^w\left(B,\tilde{B},G^s,G^x\right)+X^w\left(B,\tilde{B},G^s,G^x\right)\right).$$

The following functions determine the allocation:

$$\begin{aligned}
C^s\left(B\right) &\equiv (1-\beta)B\\
E^w\left(B,G^s,G^x\right) &\equiv 1-E^s\left(B\right)-G^x-G^s\\
S^w\left(B,\tilde{B},G^s,G^x\right) &\equiv \min\left\{\max\left\{\tilde{B}-B,0\right\},E^w\right\}\\
X^w\left(B,\tilde{B},G^s,G^x\right) &\equiv \frac{E^w-\min\left\{\max\left\{\tilde{B}-B,0\right\},E^w\right\}}{q}\\
q\left(B,\tilde{B},G^s,G^x\right) &\equiv \mathcal{A}^{-1}\left(\mu\right)\\
\mu\left(B,\tilde{B},G^s,G^x\right) &\equiv G^x+X^w q.
\end{aligned}$$

The worker's total consumption is:

$$C^w=S^w\left(B,\tilde{B},G^s,G^x\right)+X^w\left(B,\tilde{B},G^s,G^x\right).$$

Because

$$\mathcal{Y}=C^w+C^s+G^s+\frac{G^x}{q\left(\mu\right)},$$

but $C^s\left(B\right)$ is independent of the government's expenditure, we obtain: $d\mathcal{Y}=dC^w+dG^s+d\left(G^x/q\left(\mu\right)\right)$.

Thus, we have that the government's expenditure multiplier for expenditure of type $i=x,s$ relates to the worker's consumption as follows:

$$\mathcal{M}^s\left(B,\tilde{B}\right)\equiv\frac{d\mathcal{Y}}{dG^s}=\frac{dC^w}{dG^s}+1$$

and

$$\begin{aligned}
\mathcal{M}^x\left(B,\tilde{B}\right) &\equiv \frac{d\mathcal{Y}}{dG^x}=\frac{dC^w}{dG^x}+\frac{1}{q}-\frac{G^x}{q^2}\frac{dq}{d\mu}\frac{d\mu}{dG^x}\\
&= \frac{dC^w}{dG^x}+\mathcal{A}\left(\mu\right).
\end{aligned}$$

where the last equation is because we are interested in obtaining the government's infinitesimal multiplier, evaluated at $G^x = G^s = 0$. We have defined these multipliers as:

$$\mathcal{M}^i\left(B, \tilde{B}\right) \equiv \left.\frac{d\mathcal{Y}}{dG^i}\right|_{G^x=G^s=0},$$

and relating to the indirect social utility function, the change in the objective is:

$$\frac{\theta}{C^w}\frac{dC^w}{dG^i},$$

for expenditure i.

Case 1. All spot consumption $B < B^\star$. If there is only spot consumption:

$$\begin{aligned}
C^s(B) &\equiv (1-\beta)B \\
C^w\left(B, \tilde{B}, G^s, G^x\right) &\equiv 1 - (1-\beta)B - G^x - G^s \\
X^w\left(B, \tilde{B}, G^s, G^x\right) &\equiv 0 \\
q\left(B, \tilde{B}, G^s, G^x\right) &\equiv \mathcal{A}^{-1}(\mu) \\
\mu\left(B, \tilde{B}, G^s, G^x\right) &\equiv G^x.
\end{aligned}$$

Since worker consumption is independent of q, we have that:

$$\frac{dC^w}{dG^i} = -1$$

for both $i \in \{x, s\}$. In turn, we have that

$$\mathcal{M}^s\left(B, \tilde{B}\right) = \frac{dC^w}{dG^s} + 1 = 0.$$

Likewise, for chained expenditures, we have:

$$\mathcal{M}^x\left(B, \tilde{B}\right) = \frac{dC^w}{dG^x} + \mathcal{A}(0^+) = -1 + \mathcal{A}(0^+) < 0.$$

For the random and efficient assignments, $\mathcal{A}(0^+) = \delta$, so this multiplier equals $\delta - 1$.

Case 2. Some chained consumption. If there are some chained expenditures:

$$
\begin{aligned}
C^s\left(B\right) &\equiv \left(1-\beta\right)B \\
E^w\left(B, G^s, G^x\right) &\equiv 1-\left(1-\beta\right)B - G^x - G^s \\
C^w\left(B, \tilde{B}, G^s, G^x\right) &\equiv S^w + X^w \\
S^w\left(B, \tilde{B}, G^s, G^x\right) &\equiv \max\left\{\tilde{B} - B, 0\right\} \\
X^w\left(B, \tilde{B}, G^s, G^x\right) &\equiv \frac{E^w - \max\left\{\tilde{B} - B, 0\right\}}{q} \\
q\left(B, \tilde{B}, G^s, G^x\right) &\equiv \mathcal{A}^{-1}\left(\mu\right) \\
\mu\left(B, \tilde{B}, G^s, G^x\right) &\equiv G^x + X^w q.
\end{aligned}
$$

Rewriting the last three identities using $\mathcal{A}\left(\mu\right)$ instead of q we have

$$
\begin{aligned}
X^w &\equiv \mathcal{A}\left(\mu\right)\left(E^w - \max\left\{\tilde{B} - B, 0\right\}\right) \\
\mu &\equiv G^x + E^w - \max\left\{\tilde{B} - B, 0\right\}.
\end{aligned}
$$

Substituting E^w we have, naturally,

$$
\mu = 1 - \underbrace{\left(\left(1-\beta\right)B + G^s + \max\left\{\tilde{B} - B, 0\right\}\right)}_{\text{spot exp.}}. \tag{90}
$$

From here, we obtain that:

$$
\frac{dX^w}{dG^x} = -\mathcal{A}\left(\mu\right) + \mathcal{A}'\left(\mu\right)\frac{d\mu}{dG^x}\left(E^w - \max\left\{\tilde{B} - B, 0\right\}\right).
$$

Since $\frac{d\mu}{dG^x} = \frac{dS^w}{dG^x} = 0$, we have that:

$$
\frac{dC^w}{dG^x} = -\mathcal{A}\left(\mu\right).
$$

Hence, the government multiplier for chained expenditures is:

$$
\mathcal{M}^x\left(B, \tilde{B}\right) = \frac{dC^w}{dG^x} + \mathcal{A}(\mu) = -\mathcal{A}(\mu) + \mathcal{A}(\mu) = 0.
$$

Next, consider a change in spot government expenditures, holding $G^x = 0$. Observe that:

$$X^w \equiv \mathcal{A}(\mu)\left(1-(1-\beta)B - G^x - G^s - \max\left\{\tilde{B}-B,0\right\}\right).$$

Hence,

$$\begin{aligned} dX^w &= -\mathcal{A}(\mu)\,dG^s + \mathcal{A}'(\mu)\left(\frac{X^w}{\mathcal{A}(\mu)}\right)d\mu \\ &= -\mathcal{A}(\mu)\,dG^s + \mathcal{A}(\mu)\,\epsilon_\mu^{\mathcal{A}}\left(\frac{X^w}{\mathcal{A}(\mu)}/\mu\right)d\mu. \end{aligned}$$

The second line follows from:

$$\mathcal{A}(\mu)\,\mu \equiv \mathcal{A}(\mu)\,G^x + X^w.$$

Also, following this condition, we have that:

$$\mathcal{A}(\mu)\left(1+\epsilon_\mu^{\mathcal{A}}\right)d\mu \equiv dX^w.$$

Combining the differentials evaluated at $G^x = 0$, we obtain:

$$d\mu = -dG^s.$$

Hence,

$$\frac{dC^w}{dG^s} = \frac{dX^w}{dG^s} = -\mathcal{A}(\mu)\left(1+\epsilon_\mu^{\mathcal{A}}\right).$$

Following the relationship with the fiscal multiplier, we obtain:

$$\mathcal{M}^s\left(B,\tilde{B}\right) = 1 - \mathcal{A}(\mu) - \mathcal{A}(\mu)\,\epsilon_\mu^{\mathcal{A}}.$$

Summary. We summarize the results:

$$\frac{dC^w}{dG^x} = \begin{cases} -1 & B < B^\star\left(\tilde{B}\right) \\ -\mathcal{A}(\mu) & B > B^\star\left(\tilde{B}\right) \end{cases}$$

and

$$\frac{dC^w}{dG^s} = \begin{cases} -1 & B < B^\star\left(\tilde{B}\right) \\ \\ -\mathcal{A}\left(\mu\right)\left(1+\epsilon_\mu^{\mathcal{A}}\right) & B > B^\star\left(\tilde{B}\right). \end{cases}$$

Finally, recall that the government multipliers relate to the change in consumption as follows:

$$\mathcal{M}^s\left(B,\tilde{B}\right) \equiv \frac{d\mathcal{Y}}{dG^s} = \frac{dC^w}{dG^s} + 1$$

and

$$\mathcal{M}^x\left(B,\tilde{B}\right) \equiv \frac{dC^w}{dG^x} + \mathcal{A}\left(\mu\right)$$

with:

$$\frac{d\mu}{dG^x} = \begin{cases} 1 & B < B^\star\left(\tilde{B}\right) \\ \\ 0 & B > B^\star\left(\tilde{B}\right). \end{cases}$$

Therefore, adding terms:

$$\mathcal{M}^x\left(B,\tilde{B}\right) = \begin{cases} \mathcal{A}(0^+) - 1 & B < B^\star\left(\tilde{B}\right) \\ \\ 0 & B > B^\star\left(\tilde{B}\right). \end{cases}$$

Hence,

$$\mathcal{M}^s\left(B,\tilde{B}\right) = \begin{cases} 0 & B < B^\star\left(\tilde{B}\right) \\ \\ 1 - \mathcal{A}\left(\mu\right)\left(1+\epsilon_\mu^{\mathcal{A}}\right) & B > B^\star\left(\tilde{B}\right). \end{cases}$$