The B.E. Journal of Macroeconomics

Contributions

Volume 9, *Issue* 1 2009 *Article* 17

Settlement Systems

Benjamin Lester*

*University of Western Ontario, blester@uwo.ca

Recommended Citation
Benjamin Lester (2009) "Settlement Systems," *The B.E. Journal of Macroeconomics*: Vol. 9: Iss. 1 (Contributions), Article 17.

Settlement Systems*

Benjamin Lester

Abstract

We construct a general equilibrium model in which credit is used as a medium of exchange, and banks participate in a settlement system to finalize their customers' transactions. We study the optimal settlement system design, and find that a trade-off arises endogenously within the model. A higher frequency of settlement and more costly intra-day borrowing policies limit the banks' accumulation of liabilities and promote conservative reserve management, hence limiting the risk of default should a bank fail for any reason. However, such policies also imply higher banking costs, less interest paid on deposits, and less capital allocated by banks to productive investments. After characterizing equilibrium and welfare, we parameterize the economy and analyze how the optimal settlement system policies depend on several features of the economy, including the risk of bank failure, the fragility of the settlement system, the volume of trade by banks' customers, and the rate of return on investments available to banks.

KEYWORDS: payment systems, liquidity, banking

*I am particularly indebted to Randy Wright for his guidance throughout the writing of this paper. I thank Aleksander Berentsen, Ed Green, Antoine Martin, Stephen Millard, Ed Nosal, David Mills, Deniz Selman, Ludo Visschers, Matthew Willison, two anonymous referees, the editor, and seminar participants at the Federal Reserve Bank of Cleveland, the Federal Reserve Bank of New York, the Bank of Canada, the SED in Vancouver, and the Penn Macro Lunch for helpful commentary. All errors are my own.

1 Introduction

Many purchases, particularly of large value, are paid for with *credit* in the following sense: a buyer transfers to the seller a claim on funds held at his financial institution, which the seller can subsequently redeem at her financial institution. In such a transaction, the exchange of goods and the finalization of payment are separate events, and therefore require a specification for when and how the actual funds are delivered from the buyer's financial institution to the seller's. Such a specification, called a *settlement system*, is described by Zhou (2000) as a "contractual and operational arrangement that banks and other financial institutions use to transfer funds to each other."

As a result of recent technological innovations and increased financial integration, both the regularity and the value of these transactions have increased dramatically. For example, in the United States the average daily value of transfers processed by the two primary settlement systems exceeds three trillion dollars, or approximately one third of annual GDP.[1] Since the ability and confidence of market participants to make and receive payments is crucial to both commercial and financial trade, it follows that an efficient, stable settlement system is an essential component of a well-functioning modern economy.

A settlement system is characterized by two key specifications. The first is the frequency with which banks must settle their outstanding liabilities with other banks. The two most common specifications are a *deferred* settlement system, in which liabilities are settled on a net basis at the end of the trading day, and a *real-time* settlement system, in which liabilities are continuously settled. The second key specification is at what price (if at all) the operator of the settlement system is willing to extend intra-day credit to those banks that do not have sufficient liquid assets to pay their outstanding liabilities at the time of settlement.

A trade-off between cost and risk arises in the determination of the optimal settlement frequency and intra-day borrowing policy. First, increasing the frequency of settlement implies more transactions must be sent and received, and thus an increase in the costs associated with collecting, transferring, and monitoring payments from one financial institution to another. These costs are often referred to as *resource* costs. Secondly, increasing the frequency of settlement or the cost of intra-day borrowing implies that banks allocate a smaller fraction of deposits into productive yet illiquid investments. These

[1] The two primary settlement systems in the United States are Fedwire and CHIPS. It should be noted that these statistics are similar across most developed countries; for example, in the United Kingdom and Canada the average daily value of transactions is approximately 20% and 15%, respectively, of annual GDP.

opportunity costs of foregone revenue from investment are often referred to as *financial* costs. In a riskless world with perfect commitment between banks, the optimal settlement system would simply be the least costly, and therefore frequent settlement and costly intra-day liquidity would be undesirable characteristics.

However, in reality there *is* risk that a financial institution will be unable to deliver funds to those banks awaiting payment, and therefore must default on its outstanding liabilities; indeed, recent events in the financial sector have brought such risks to the forefront of policymakers' concerns. The settlement system in place will be a crucial determinant of how the economy as a whole is affected by such risks.[2] Increased settlement frequency and more costly intra-day borrowing policies mitigate this risk in two important ways. First, as the frequency of settlement increases, the number of parties awaiting payment decreases. Therefore, a settlement system that reduces the time between the initial transaction and the finalization of payment decreases the potential size of defaults between banks. Secondly, since banks retain a higher level of reserves under systems with more frequent settlement or more costly borrowing policies, they are more equipped to honor outstanding liabilities in the case of insolvency.

The proper choice of settlement frequency and intra-day borrowing costs have recently become the subject of significant debate, as central banks world-wide have re-examined their policies. For example, over the past two decades many European countries concluded that the risks posed by deferred settlement systems were too large, and thus pushed for the implementation of high frequency real-time settlement systems.[3] However, there are those who argue that the merits of real-time settlement systems did not outweigh the costs.[4] In the U.S., the settlement system operated by the Federal Reserve (Fedwire) has recently changed its intra-day borrowing policy; liquidity was provided at zero cost until 1994, at which time a fee of twenty-four basis points was introduced. Fedwire has since raised this fee to curb excessive intra-day borrowing, and continues to carefully monitor this policy's effect on cost and risk.

[2]There are several papers that assess the performance of a settlement system after a large disruption to the financial sector. See Bernanke (1990) for analysis of clearing and settlement during the 1987 stock market crash, and McAndrews and Potter (2002) and Fleming and Garbade (2002) for an account of settlement performance after the September 11th terrorist attacks.

[3]As Bech and Hobijn (2007) report, the number of central banks using a real-time settlement system grew from three in 1985 to ninety in 2005. In particular, Germany implemented a real-time system in 1987, Italy in 1989, the United Kingdom in 1996, and both France and Spain in 1997. See Martin (2005) for more discussion.

[4]See, for example, Selgin (2004).

Though the determination of the optimal settlement system is an important and pressing issue, it has received surprisingly little attention in the academic literature. The current paper is a response to this observation, and we proceed in three steps. First, we introduce a general equilibrium model with two sectors. On the one hand, there is a banking (or financial) sector, in which banks accept deposits from their customers, allocate a fraction of these deposits to productive investments, and retain the remainder in order to settle their customers' transactions through a *settlement system*, which is assumed to be operated by the central bank. On the other hand, there is a trading sector, in which each bank's customers randomly meet and trade with other banks' customers in an anonymous, decentralized market, using credit as a medium of exchange. To start, we assume that there is no risk in the banking sector. This simple framework provides us with a clean, general equilibrium model in which a relatively serious model of banks is integrated into a setting with real transactions that has a micro-founded role for bank liabilities as a media of exchange.[5] As a result, decisions in the banking sector affect the behavior of agents in the trading sector, and vice versa.

Second, we introduce risk into the financial sector by assuming that a stochastic fraction of banks may suddenly fail in any given period, and we model explicitly how such failures affect the economy. In this setting, we show formally how a trade-off between cost and risk arises endogenously in the determination of the optimal settlement frequency and intra-day borrowing policy. The result is a model that allows for the characterization of the optimal settlement system design, taking into account the risk of bank failures, the opportunity costs of liquidity provision, and the real costs of a large-scale disruption to the payment system. In the third and final step of our analysis, we parameterize the model and characterize the optimal *joint* policy (i.e. the optimal combination of settlement frequency and borrowing policy) as a function of the economy's fundamentals. This exercise is not a calibration per se, but rather used to illustrate the main insights provided by the model, which are as follows.

[5]There are several reasons why this random matching framework is appealing for the current analysis. First, *some* medium of exchange is essential, and thus some form of settlement finality is also essential. Moreover, the use of *credit* as a medium of exchange is micro-founded, in the sense that agents choose whether or not to use the banking system given the costs and risks. This feature provides some discipline as we consider various policy choices, as all equilibria must satisfy the incentive constraints of buyers and sellers. Finally, the model assumes that bilateral transactions occur randomly across time, which is consistent with the stochastic evolution of intra-day net liabilities introduced below.

We find that the optimal intra-day lending policy tends to be more lenient under systems with more frequent settlement. There are two reasons for this: first, banks have greater liquidity needs under high frequency settlement systems, and thus reserve management is more sensitive to the cost of credit; and second, high frequency settlement reduces risk independently of banks' reserves, thereby reducing the need for the central bank to curb risk through the intra-day credit rate.

We also find that a high frequency settlement system is optimal when the economy is more vulnerable to bank failure or when systemic risk (the possibility that defaults lead to a disruption in the settlement system) is high. Moreover, the model uncovers an interesting relationship between the optimal settlement system, welfare, and the volume of trade in the real sector. More trade implies more consumption, but it also leads to greater volatility of banks' intra-day liabilities. Such volatility entails higher banking costs and more potential for default. As a result, we find that higher frequency settlement systems are preferred when the level of real economic activity is high.

1.1 Related Literature

Much of the previous work on settlement systems, beginning with Freeman (1996), utilizes relatively abstract models to determine the optimal intra-day borrowing policy. In these models, the agents *are* the banks, and the need for a third party to provide liquidity arises from a carefully specified spatial and intertemporal lack of coordination. The first generation of these papers, which would include Freeman (1999) and Zhou (2000), introduced default as an exogenous event and concluded that intra-day liquidity should be provided free of charge. More recent work by Martin (2004), Mills (2006) and Koeppl et al. (2004) has endogenized risk by modeling default as the product of a moral hazard. In this context, they find that the free provision of intra-day liquidity may not be optimal.

The approach in these papers is appealing because there is a natural definition of welfare and a transparent role for a welfare-improving agent that provides the market with liquidity. However, there are two limitations. First, since many of these models require a highly stylized timing sequence, they do not easily accommodate the analysis of settlement frequency.[6] Therefore, they cannot consider settlement design and liquidity provision as a single policy question. Secondly, by modeling banks and consumers/producers as a single

[6]See Koeppl et al. (2008) for the most recent attempt to incorporate settlement frequency into this type of environment.

agent, these authors abstract from issues of investment, real production, and the ultimate effects of settlement system design on consumption. Therefore, they cannot address one of the motivating concerns in this area of research: that a disruption to the settlement system would have adverse effects for real economic activity.

Kahn and Roberts (1998) overcome the first of these limitations by developing a model with a more realistic notion of banks and a timing sequence that allows for the analysis of settlement frequency. In this model, banks must choose the proper portfolio of liquid and non-liquid assets given that intra-day liabilities are stochastic. In the presence of a moral hazard problem similar to that discussed above, the model successfully generates the trade-off between cost and risk that is inherent to the choice of settlement frequency. However, as they ignore the behavior of the agents involved in the underlying transactions, Kahn and Roberts (1998) require arbitrary weights to conduct welfare and policy analysis.[7]

2 Benchmark Model: No Risk

Consider a discrete time, infinite horizon model in which each period, or "day", consists of a finite number of sub-periods. There is a mass B of ex-ante identical banks in the banking sector, and a mass of agents in the trading sector, which we normalize to 1. Both banks and agents in the trading sector can produce a general good. The cost of producing y units of the general good is $C(y) = y$, and yields utility $U(y) = y$.

A fraction K of agents in the trading sector possess a single, indivisible unit of capital and the remainder do not. Those agents with capital will be referred to as *buyers* and those without capital as *sellers*. Capital cannot be consumed. As we discuss in detail below, banks have access to a technology that allows them to use capital to produce the general consumption good, but agents in the trading sector do not have access to this technology.[8] Instead, buyers can use capital as a medium of exchange, or they can deposit it at a bank.

[7]One issue that the current paper does not address is the strategic timing of payment finalization and the risk of so-called "gridlock."See Martin and McAndrews (2008), Kahn et al. (2003), and Martin and McAndrews (2008) for analysis of this issue.

[8]There are a variety of reasons why banks have access to certain investment opportunities and technologies that individuals do not: they tend to have a comparative advantage in finding, managing, and monitoring new projects, as well as designing and enforcing the terms of the investment. Though this is taken as given here, one could easily generate this outcome endogenously in a number of ways.

The buyers' decision to deposit capital at a bank is made at the beginning of each day. We assume that those buyers that deposit their capital do so randomly across banks, so that each bank begins the day with the same level of deposits, denoted by d. Upon receiving a deposit, a bank issues the depositor some form of receipt that can then be used to purchase goods in the trading sector. Let us loosely refer to this receipt as a "check."Banks make a flow interest payment on each deposit, ϕ, which is an endogenous variable to be determined in equilibrium. We assume that ϕ is paid at the time of deposit, and is paid in the form of the general good.

In what follows, we begin by describing in detail the behavior of banks, taking as given the behavior of agents in the trading sector. Then we will consider the problem of buyers and sellers, given the behavior of banks, before defining a general equilibrium.

2.1 The Banking Sector

Before the day's trading commences, banks choose a fraction α of initial deposits to retain as reserves, and invest the remainder $i \equiv (1-\alpha)d$. Banks have access to the following technology: a bank that invests i units of capital will receive $f(i)$ units of the consumption good at the end of the period, as well as recovering the initial i units of capital, where f is a continuous, strictly concave function of i.[9]

In the trading sector, each bank has a mass d of active buyers. Each seller is pre-assigned a bank to deposit any payments received, and we assume that sellers are uniformly distributed across banks. In a transaction in which a buyer gives the seller a check, we assume that the seller can deposit the check immediately. Therefore, after such a transaction occurs, the buyer's bank owes the seller's bank one unit of capital. More generally, throughout the day, a bank accrues *due-to*'s and *due-from*'s as its buyers and sellers, respectively, trade with other banks' customers in the marketplace. A bank's net liabilities to other banks at any time are equal to the difference between due-to's and due-from's.[10]

Suppose that transactions occur at T discrete intervals throughout each day, and that a bank has cumulative intra-day net liabilities $l_t \in [\underline{L}_t, \bar{L}_t]$ at each $t \in \mathcal{T} \equiv \{1, 2, ..., T\}$, where $\bar{L}_t$ ($\underline{L}_t$) is the maximum (minimum) net

[9] The banks' technology allows them to make capital divisible.

[10] In what follows, we will often use the phrase "net liabilities"to refer to the amount that a bank owes (or is owed) by other banks in the settlement system. This should not be confused with its total net obligations, which would naturally include, for example, its liabilities to its depositors.

position that occurs at time t with strictly positive probability, and $-d \leq \underline{L}_t < 0 < \bar{L}_t \leq d$ for all $t \in \mathcal{T}$. Denote the maximum possible net liability position $\bar{L} = \max_{t \in \mathcal{T}}\{\bar{L}_1, ..., \bar{L}_T\}$. Therefore, a bank's intra-day net liabilities is a T x 1 random variable $l \equiv (l_1, ..., l_T)$ with density $\lambda(\cdot)$.[11] Any outstanding liabilities are settled at the end of each day, so that each period begins with $l_0 = 0$. Moreover, each night deposits are reallocated across banks so that the following morning all banks begin with d deposits.

2.1.1 Notation and Key Concepts

One key attribute of a settlement system is the frequency with which banks must settle net liabilities incurred since the previous settlement. Let $\mathcal{S} \subset \mathcal{T}$ denote the subset of points in time in which banks settle. Under this formulation, a slightly narrow definition of *increased settlement frequency* is required: for any settlement system which specifies a subset of intervals at which settlement occurs, let us say that increasing the settlement frequency implies adding at least one additional interval to the existing subset. Formally, we say that $\mathcal{S}'' \subset \mathcal{T}$ has a higher frequency of settlement than $\mathcal{S}' \subset \mathcal{T}$ if and only if $\mathcal{S}'$ is a strict subset of $\mathcal{S}''$. Also, for any given point in time t, let s_t denote the time of the previous settlement. That is, let $s_t = \max\{t' \in \mathcal{S} : t' < t\}$ for $t > 1$, with $s_1 = 0$.

The settlement system is operated by a third party which we refer to as the "central bank", as this is both convenient and consistent with many settlement systems world-wide. Each bank has an account at the central bank, where it keeps its reserves. At the time of settlement, the central bank will transfer the appropriate amounts of capital from the accounts of banks that have accrued positive net liabilities since the previous settlement to the accounts of banks that have accrued negative net liabilities since the previous settlement. Therefore, as the operator of the settlement system, the central bank must collect, monitor, and execute payment finalization for its member banks. The costs of such activity are referred to as *resource costs*, and they can be substantial.[12] Therefore, suppose that there is a cost h that the central

[11] One might think, at first, that if each bank has a continuum of buyers and sellers, then $l_t = 0$ at all t by the law of large numbers. However, as we illustrate later, we are making no assumptions on the independence of meetings between buyers and sellers; the probability that buyers from one bank meet sellers from another, and the time of these meetings, may very well be correlated. Therefore, $\lambda(\cdot)$ will not, in general, be a degenerate distribution.

[12] For example, Humphrey et al. (1996) estimate that the real resource costs of operating all payments systems – of which the primary settlement systems are a large portion – is approximately 1% - 1.5% of annual GDP. See Berger et al. (1996) for a more thorough discussion of resource costs.

bank incurs each time it must transfer capital to or from a bank. The following definition characterizes the resource costs of the central bank for a given day.

Definition 1. *Let the central bank's* ***resource costs under settlement frequency*** $\mathcal{S}$ *be denoted* $H(\mathcal{S})$*, and be defined:*

$$H(\mathcal{S}) = B\left[\sum_{t\in\mathcal{S}} \int h \cdot \mathbf{1}_{\{l_t \neq l_{s_t}\}} \lambda(l) dl\right]. \tag{1}$$

Note that $\mathbf{1}_{\{l_t \neq l_{s_t}\}}$ represents the indicator function which assumes the value 1 if a given bank's net liabilities at the current settlement period (t) are not equal to its net liabilities at the previous settlement period (s_t). In the following definition, we characterize the amount of reserves that a bank holds in its account at the central bank at any given time, which will be a crucial statistic in our analysis.

Definition 2. *Let a bank's* ***reserves at time*** t *for any realization of net liabilities* l*, reserve ratio* α*, and settlement frequency* $\mathcal{S}$ *be denoted* $R_t(l,\alpha;\mathcal{S})$*, and be recursively defined:*

$$R_t(l,\alpha;\mathcal{S}) = \begin{cases} \alpha d & \text{if } t = 0 \\ \max\{0, R_{s_t}(l,\alpha;\mathcal{S}) - (l_t - l_{s_t})\} & \text{if } t \in \mathcal{S} \\ R_{s_t}(l,\alpha;\mathcal{S}) & \text{if } t \notin \mathcal{S}. \end{cases}$$

As the formula suggests, a bank's account balance at the central bank is non-negative, increases at settlements in which the bank receives payment (when $l_t - l_{s_t} < 0$ for $t \in \mathcal{S}$), and decreases when a bank must settle outstanding liabilities (when $l_t - l_{s_t} > 0$ for $t \in \mathcal{S}$). Notice that for any t that does not correspond to a settlement period, the level of reserves does not change. Should a bank's obligations exceed its reserves at any settlement time $t \in \mathcal{S}$, it must borrow the difference, $(l_t - l_{s_t}) - R_{s_t}(l,\alpha;\mathcal{S})$, from the central bank.[13]

Definition 3. *Let a bank's* ***total intra-day borrowing requirement*** *for a given realization of net liabilities* l*, reserve ratio* α*, and settlement frequency* $\mathcal{S}$ *be denoted* $b(l,\alpha;\mathcal{S})$ *and be defined:*

$$b(l,\alpha;\mathcal{S}) = \sum_{t\in\mathcal{S}} \max\{0, (l_t - l_{s_t}) - R_{s_t}(l,\alpha;\mathcal{S})\}.$$

[13]Note that, under the definition below, a bank's total intra-day borrowing requirement is simply the maximum negative balance it achieves during the day; this is derived explicitly in equation (36). Thus we are abstracting from how many sub-periods a bank maintains a particular balance during the day, and instead looking at only the total amount a bank would have to borrow in order to settle its liabilities at each settlement period within a day.

In addition to the frequency of settlement, the conditions under which a central bank provides intra-day liquidity is the second crucial characteristic of a settlement system. We assume that the central bank sets a fixed rate $k \geq 0$ at which banks can borrow throughout the day.[14] Therefore, the policy analysis herein focuses on two features of a settlement system: the frequency of settlement ($\mathcal{S}$) and the terms of intra-day borrowing (k). Let $\mathcal{P} \equiv (\mathcal{S}, k)$ denote a given policy. Since banks choose their reserve levels before the start of the day, they must project borrowing fees for any given level of reserves.

Definition 4. *Let a bank's* ***expected borrowing fees****, given reserve ratio* α *and policy* $\mathcal{P}$*, be denoted* $\kappa(\alpha; \mathcal{P})$ *and be defined:*

$$\kappa(\alpha; \mathcal{P}) = \mathbb{E}_l \Big[k \cdot b(l, \alpha; \mathcal{S}) \Big] = k \int b(l, \alpha; \mathcal{S}) \lambda(l) dl.$$

It is assumed that borrowing fees are paid in the form of the general good. We assume that the borrowing fees that are collected by the central bank, net of resource costs, are distributed at the end of each period to agents in the trading sector in the form of a lump sum transfer. We denote the value of this lump-sum transfer for a particular policy $\mathcal{P}$ by

$$\tau(\mathcal{P}) = \kappa \left[\alpha(\mathcal{P}); \mathcal{P} \right] B - H(\mathcal{S}). \tag{2}$$

2.2 The Problem of the Bank

A bank generates revenue, in the form of the general good, by investing a portion of its deposits. It pays interest, ϕ, to each depositor, and expects to pay $\kappa(\alpha; \mathcal{P})$ in borrowing fees from the central bank. Therefore, a bank's expected profits are given by

$$\pi(\phi, \alpha; \mathcal{P}) = f[(1 - \alpha)d] - \kappa(\alpha; \mathcal{P}) - \phi d. \tag{3}$$

Given ϕ, the optimal level of reserves that a profit-maximizing bank will choose under settlement system $\mathcal{P}$, which we denote $\alpha(\mathcal{P})$, is the solution to the first order condition

$$-\kappa'(\alpha; \mathcal{P}) = d[f'(i)]. \tag{4}$$

[14] Such a formulation is consistent with current practice in the U.S. settlement system, Fedwire, and much of the literature (see Martin and McAndrews (2008)). There are, of course, alternative arrangements such as collateral requirements and borrowing caps, which are utilized in other settlement systems. The manner in which borrowing fees are assessed is an interesting issue in and of itself (see Mills (2006)). The current model could be generalized to consider arbitrary borrowing fee schedules (i.e. arbitrary functions of b). However, for both simplicity and tractability, we focus here on the case of a constant fee schedule.

This first order condition has the usual interpretation: the marginal benefit from an increase in the reserve ratio, resulting from a reduction in expected borrowing fees, is equated to the marginal cost of foregone revenue from investing. Moreover, in a perfectly competitive banking sector, the zero expected profit condition implies

$$\phi = \frac{1}{d}\Big[f[(1-\alpha)d] - \kappa(\alpha;\mathcal{P})\Big]. \tag{5}$$

Therefore, given $\alpha(\mathcal{P})$, the interest rate on deposits $\phi(\mathcal{P})$ follows immediately from (5). Moreover, the aggregate level of production given policy $\mathcal{P}$ is given by

$$Y(\mathcal{P}) = f\left\{[1-\alpha(\mathcal{P})]\, d\right\} B. \tag{6}$$

2.3 The Trading Sector

The trading sector is modeled in the spirit of He et al. (2005). There is a $[0, 1]$ continuum of agents and, as mentioned earlier, a fraction K are initially endowed with one unit of capital. Capital is indivisible and agents hold either zero or one unit; an agent with capital is a buyer, and an agent without capital is a seller. Before entering the market, a buyer has the option to deposit his capital at the bank. In exchange, he receives some form of receipt (a "check") and an interest payment ϕ, which comes in the form of the general good. Let us denote the probability that a representative buyer deposits his money at the bank (the buyer's strategy) as θ.

We study equilibria in which both capital and checks are acceptable as media of exchange.[15] Therefore, in a model without risk, the buyer's strategy is trivial: he strictly prefers to deposit his capital if $\phi > 0$, he strictly prefers not to deposit if $\phi < 0$ (i.e. if the bank charges a fee), and he is indifferent if $\phi = 0$.[16] In the next section, when we introduce risk, this decision will no longer be trivial. For now, we concentrate on equilibria with $\phi \geq 0$ and $\theta = 1$.

Agents meet stochastically in an anonymous, bilateral matching market. Each agent is matched with another agent once each day. With probability K that agent is a buyer, and with probability $1 - K$ that agent is a seller. The

[15] It is by now well understood that in the trading environment described here, *some* medium of exchange is essential; see, for example, Kiyotaki and Wright (1993), Kocherlakota (1998), and Wallace (2001). It is also well known that inside money can play this role, so long as agents cannot produce it themselves; see, for example, Cavalcanti and Wallace (1999a) and Cavalcanti and Wallace (1999b). Below, we state explicitly the incentive compatibility constraints that support this type of equilibrium.

[16] We will assume that he deposits his capital if he is indifferent.

sub-period at which this meeting occurs ($t \in \{1, 2, ..., T\}$), and the bank that this agent is associated with, is not necessarily random; for different matching processes, a different distribution $\lambda(\cdot)$ will be generated for banks. When a buyer meets a seller, there is probability x that he likes the specialized good that the seller can produce. In this case, the seller produces the good at cost c in exchange for the buyer's check, and the buyer immediately consumes the good and receives utility u. The probability of a double-coincidence of wants is zero. Goods are indivisible, non-storable, and they are produced, traded, and consumed immediately. After an agent sells her good, she leaves the market and deposits the check at her bank. When the bank receives payment finalization, her account is credited one unit and she becomes a buyer the following day.

2.4 Value Functions and Incentive Constraints

Let us denote the value function of a seller by V_0, and the value function of a buyer with one unit of capital deposited at the bank by V_1, so that

$$V_0 = Kx(-c + \beta V_1) + (1 - Kx)\beta V_0 + \tau \tag{7}$$
$$V_1 = (1 - K)x(u + \beta V_0) + [1 - (1 - K)x]\beta V_1 + \phi + \tau. \tag{8}$$

Consider V_0, the value function of a seller. With probability K the seller meets a buyer, and with probability x the buyer likes his good and trade ensues: he incurs production cost c and receives payment. In addition, all agents receive the lump sum transfer τ. Similar reasoning applies to the buyer's value function. An equilibrium in which agents use banks must satisfy three incentive compatibility conditions:

$$\beta(V_1 - V_0) \geq c \tag{9}$$
$$u \geq \beta(V_1 - V_0) \tag{10}$$
$$\phi \geq 0. \tag{11}$$

The first two conditions ensure that both producers and consumers, respectively, want to trade in a bilateral match. The final condition ensures that agents with money choose to use banks. Given $\phi \geq 0$, one can show that

$$\beta(V_1 - V_0) = \frac{[\phi + Kxc + (1 - K)xu]}{r + x},$$

and therefore, letting r be the time discount factor so that $\beta = 1/(1+r)$, (9) and (10) simplify to

$$c \leq \frac{\phi + (1-K)xu}{r + (1-K)x} \tag{12}$$

$$u \geq \frac{\phi + Kxc}{r + Kx}. \tag{13}$$

The first condition ensures that the gains from trade – including the interest paid on future deposits – is sufficiently large that a seller is willing to incur the costs of production in order to become a buyer. The second condition ensures that the interest earned on deposits is not so high that buyers are unwilling to withdraw their deposits from the bank.

2.5 Equilibrium and Welfare

We now define an equilibrium in which all buyers utilize the banking system.

Definition 5. *Given settlement policy $\mathcal{P}$, a* ***banking equilibrium*** *is a pair of strategies (α^*, θ^*) with $\theta^* = 1$ such that*

- *given $\mathcal{P}$, α^* satisfies the bank's optimality condition in (4). Moreover, α^* uniquely determines interest rates ϕ^* according to (5) and transfers τ^* according to (2).*
- *given ϕ^* and τ^*, the constraints (11) - (13) are satisfied.*
- *bank deposits equal the supply of capital that agents deposit, so that $d = K/B$.*

Recall that banks make zero profits, so that welfare is simply a weighted average of the expected lifetime utility of buyers and sellers. In particular,

$$W = KV_1 + (1-K)V_0. \tag{14}$$

Substituting (2), (5), (7), and (8) into (14) yields

$$W \propto B\left\{f\left[(1-\alpha)\frac{K}{B}\right]\right\} - H + K(1-K)x(u-c), \tag{15}$$

where we've suppressed the arguments of α and H for convenience. Note that (15) confirms an earlier conjecture: in the absence of risk, the optimal policy $\mathcal{P} \equiv \{\mathcal{S}, k\}$ is one that maximizes investment and minimizes expenditures on resource costs. Therefore, the optimal policy is $k = 0$ (so that $\alpha(\mathcal{P}) = 0$) and $\mathcal{S} = \{T\}$ (so that settlement occurs only when necessary, which is in the last period of each day by assumption).

3 Bank Failure, Default, and Systemic Risk

In the previous section, we introduced an internally consistent general equilibrium model of exchange in which banks played three roles: accepting deposits, investing, and settling the payments of their customers. We showed how the optimal settlement system was the least costly; in a riskless world, efficiency was achieved by minimizing the number of settlement periods and maximizing the fraction of deposits that were invested. Though this model serves as a useful benchmark, it abstracts from the primary concern of policymakers: the risks posed by the failure of a major financial institution. In this section, we introduce the possibility of bank failure, and illustrate how a trade-off between efficiency and risk arises endogenously. In the presence of this risk, we show how the optimal settlement system balances this trade-off, and how the optimal policy responds to changes in the economic environment.

3.1 Bank Failure

Just prior to the last sub-period of the day, T, a fraction $1 - \sigma$ of banks "fail," where σ is a random variable drawn from a distribution $G(\sigma)$ with mean $\bar{\sigma}$ and support $[0, 1]$.[17] We take a very simplistic and stark approach to modeling bank failure: when a bank "fails," it loses the entirety of its capital investment, it receives no payoff ($f(i) = 0$), and it loses the technology that converts capital into the general good. Therefore, a failed bank essentially dies.[18] To keep the model as simple as possible, we assume that a new bank acquires a failed bank's capital investment at zero cost and begins operating the following period; this way, both the number of banks and the capital stock remain constant. Moreover, depositors are fully insured: a depositor of a failed bank is simply assigned to the new bank that acquired his capital.

---

[17]The assumption that this shock is realized just prior to T is purely for convenience. All of the ensuing results hold under the more general assumption that a fraction of banks fail at any $t \in \mathcal{T}$.

[18]This approach to modeling bank failures is highly stylized, though the notion that a major financial institution can lose the entirety of its capital investments is perhaps no longer as abstract as it once seemed. Nontheless, the intent here is not to provide a deep theory of *why* banks fail or how to prevent such failures, but rather to focus on the potential risks that bank failures pose to settlement systems, and how to mitigate such risks. The current approach to modeling bank failures, described in detail below, is meant to capture several important features common to most bank failures: (i) they are sudden, (ii) the banks do not internalize many (here, all) of the negative externalities that are borne by other market participants, and (iii) those failed banks with more capital held as reserves are better able to settle outstanding debts than those with fewer available funds.

Banks awaiting payment within the settlement system, however, are not insured. Since failed banks may have insufficient reserves to settle their liabilities, they may be unable to deliver payment to awaiting banks; that is, they may need to *default* on some of their outstanding liabilities. We now turn our attention to how such defaults could potentially cause an interruption to the settlement system, and how such an interruption within the banking sector could disrupt real economic activity in the trading sector. There are two important statistics to consider when a bank fails: its remaining capital held as reserves, R_{s_T}, and its remaining net liabilities to other banks, $l_T - l_{s_T}$. If $R_{s_T} \geq l_T - l_{s_T}$, then all payments are finalized, and any remaining reserves are distributed to the new banks that enter the following morning. If, on the other hand, a failed bank's remaining liabilities to other banks are greater than its reserves, so that $R_{s_T} < l_T - l_{s_T}$, then the bank must default on all those liabilities in excess of its current reserves. For the ensuing analysis, the following definition is convenient.

Definition 6. *Let a bank's* ***outstanding liabilities in excess of reserves at time*** *t for a given realization of net liabilities l, reserve ratio α, and settlement frequency $\mathcal{S}$ be denoted $\hat{l}_t(l,\alpha;\mathcal{S})$ and be recursively defined by $\hat{l}_t(l,\alpha;\mathcal{S}) = \max\{0, l_t - l_{s_t} - R_{s_t}(l,\alpha;\mathcal{S})\}$ for all $t \in \mathcal{T}$, with $\hat{l}_0(l,\alpha;\mathcal{S}) = 0$.*

Note that, for a bank that has not failed, this is simply the amount they need to borrow at time t. However, for a failed bank at time T, $\hat{l}_T(l,\alpha;\mathcal{S})$ is the size of its default. Assuming that l is i.i.d. across banks, the law of large numbers implies that the total value of defaults can be defined in the following manner.

Definition 7. *Let the* ***aggregate default under settlement system*** *$\mathcal{P}$ for a given realization of σ be denoted $\delta(\sigma;\mathcal{P})$, and the* ***expected aggregate default under settlement system*** *$\mathcal{P}$ be denoted $\Delta(\mathcal{P})$ and be defined, respectively, as*

$$\delta(\sigma;\mathcal{P}) = B(1-\sigma)\int \hat{l}_T(l,\alpha(\mathcal{P});\mathcal{S})d\lambda(l) \tag{16}$$

$$\Delta(\mathcal{P}) = \int \delta(\sigma;\mathcal{P})\ dG(\sigma). \tag{17}$$

The primary risk of bank default, and perhaps the foremost concern to policymakers, is that a sufficiently large default could halt the operation of the settlement system entirely, and thus shut down the markets it supports. In order to capture this concern, suppose that a settlement system is able to withstand (and immediately resolve) defaults up to size $\bar{\delta}$. However, should the amount of aggregate defaults exceed $\bar{\delta}$, the settlement system must shut down for a day in order to resolve the default.[19] During this disruption, banks cannot accept deposits or facilitate withdrawals, nor can they send or receive payment finalization. Therefore, when the settlement system shuts down, buyers who have deposited their capital face one period of autarky, and sellers can only hope to trade with buyers who did not deposit their capital.

In order to characterize the probability of a disruption to the settlement system, we must first characterize the maximum possible shock (the minimal value of σ) that the system can suffer without disruption. Denote this shock $\underline{\sigma}(\mathcal{P})$, and let it satisfy

$$\delta\left[\underline{\sigma}(\mathcal{P});\mathcal{P}\right] = \bar{\delta}. \tag{18}$$

The final definition, characterizing the risk of settlement system failure, now follows naturally.

Definition 8. *Let the* ***probability of a settlement system disruption under policy*** $\mathcal{P}$ *be denoted* $q(\mathcal{P})$ *and be defined*

$$q(\mathcal{P}) \equiv G\left[\underline{\sigma}(\mathcal{P})\right]. \tag{19}$$

3.2 The Trade-off Between Cost and Risk

In this section we illustrate the trade-off between cost and risk associated with the choice of settlement frequency and intra-day borrowing policy. The first result, relating the frequency of settlement to resource costs, is straightforward and thus we omit the proof. The remaining proofs can be found in the appendix.

Lemma 1. *Increasing the frequency of settlement increases resource costs. That is,* $\mathcal{S}' \subset \mathcal{S}'' \;\Rightarrow H(\mathcal{S}'') \geq H(\mathcal{S}')$.

We now turn to the relationship between those variables determined endogenously in the banking sector - namely the interest rate ϕ and the potential default size δ - and the policy choices $\mathcal{S}$ and k. We first illustrate how increasing the frequency of settlement or the cost of intra-day liquidity will increase the financial costs to banks, and subsequently decrease the interest earned by depositors.

[19] There are two things to note. First, we are assuming that $\bar{\delta}$ is independent of $\mathcal{S}$, which is not obviously true. An interesting extension would be to allow the fragility of the settlement system to depend on the frequency of settlement. I thank the editor for this suggestion. Second, we are assuming that the settlement system shuts down for a single period. It is, in fact, straight-forward to generalize the analysis so that the settlement system shuts down for any arbitrary number of periods.

Lemma 2. *For any settlement policy with costly intra-day liquidity, increasing the reserve ratio decreases the expected borrowing fees. That is, for any* $\mathcal{P}$ *with* $k > 0$, $\kappa'(\alpha; \mathcal{P}) \leq 0$, *with strict inequality for* $\alpha < \frac{\bar{L}}{d}$.

Let $\mathcal{P}'_{\mathcal{S}} = (\mathcal{S}', k)$ and $\mathcal{P}''_{\mathcal{S}} = (\mathcal{S}'', k)$ be such that $\mathcal{S}' \subset \mathcal{S}''$, so that $\mathcal{P}''_{\mathcal{S}}$ is a policy with greater frequency of settlement. Similarly, let $\mathcal{P}'_k = (\mathcal{S}, k')$ and $\mathcal{P}''_k = (\mathcal{S}, k'')$ be such that $k'' > k'$, so that $\mathcal{P}''_k$ is a policy with more costly intra-day liquidity.

Lemma 3. *For any level of reserves, increasing the frequency of settlement or the cost of liquidity increases the expected borrowing fees. In response, banks will choose to retain a larger fraction of deposits as reserves. That is, for* $j \in \{\mathcal{S}, k\}$, *we have that* $\kappa(\alpha; \mathcal{P}'_j) \leq \kappa(\alpha; \mathcal{P}''_j)$ *and thus* $\alpha(\mathcal{P}'_j) \leq \alpha(\mathcal{P}''_j)$.

We are now ready to formalize the first major result, relating settlement system policy to interest rates and the level of production from bank investment.

Proposition 1. *Interest rates on deposits and aggregate production from bank investment decrease with both settlement frequency and the cost of intra-day liquidity. That is, for* $j \in \{\mathcal{S}, k\}$, $\phi(\mathcal{P}'_j) \geq \phi(\mathcal{P}''_j)$ *and* $Y(\mathcal{P}'_j) \geq Y(\mathcal{P}''_j)$.

From Lemma 1 and Proposition 1, we conclude that increasing the frequency of settlement implies greater resource and financial costs, while increasing the cost of intra-day credit increases financial costs only.

Although increased settlement frequency and expensive intra-day borrowing are more costly, they also decrease the risk posed by bank failures. This reduction in risk can be decomposed into two distinct effects. The first, which we will call the *accumulation effect*, is specific to increases in the frequency of settlement. In words, more frequent settlement prohibits the accumulation of large outstanding net liability positions.

Lemma 4. *At any time* $t \in \mathcal{T}$ *and for any reserve ratio* α, *the size of outstanding liabilities in excess of reserves is always smaller under a system with greater settlement frequency. That is, for any* $\mathcal{S}' \subset \mathcal{S}''$, t, α, *and* l, $\hat{l}_t(l, \alpha; \mathcal{S}') \geq \hat{l}_t(l, \alpha; \mathcal{S}'')$.

The second way in which policy can reduce risk is by promoting more conservative reserve management. We will call this the *reserve effect.* Naturally, any policy that causes an increase in the level of reserves decreases the potential for a disruption to the settlement system. Since the optimal reserve ratio is increasing in both settlement frequency and the cost of intra-day borrowing, it follows that both of these policy levers can reduce risk via the reserve effect.

Lemma 5. *At any time $t \in \mathcal{T}$ and for any settlement frequency $\mathcal{S}$, increasing the reserve ratio decreases the size of outstanding liabilities in excess of reserves. That is, for any t, $\mathcal{S}$, and l, $\alpha_2 > \alpha_1 \;\Rightarrow\; \hat{l}_t(l, \alpha_1; \mathcal{S}) \geq \hat{l}_t(l, \alpha_2; \mathcal{S})$.*

Given the results above, we immediately arrive at the second major result, relating the settlement system in place to the size of potential defaults.

Proposition 2. *The size of potential defaults is decreasing in both the frequency of settlement and the cost of intra-day borrowing. That is, for any $\sigma \in supp(G)$ and for $j \in \{\mathcal{S}, k\}$, $\delta(\sigma; \mathcal{P}'_j) \geq \delta(\sigma; \mathcal{P}''_j)$.*

Lemma 3 and Proposition 2 imply that the potential cost of bank failures are decreasing in both the frequency of settlement and the cost of intra-day borrowing. Note that an increase in the frequency of settlement implies a reduction in the risk of an interruption to the settlement system because of both the accumulation and the reserve effects, whereas an increase in the cost of intra-day liquidity only reduces this risk via the reserve effect.

In summary, a trade-off between cost and risk arises endogenously in the model of banks developed above. More frequent settlement and more costly intra-day borrowing policies reduce the risk of a disruption to the trading sector. However, this reduction in risk comes at a cost; such policies also increase resource costs, decrease the amount of capital invested by banks, and reduce interest rates paid on deposits. Given this trade-off, the question arises: what features of an economy determine the optimal settlement system policy? We now return to the trading sector and incorporate the risks posed by bank failure. This allows, again, for a well-defined welfare function and the exploration of optimal policies.

3.3 Value Functions and Incentive Constraints

We first define the value functions when the settlement system is operational. The value function of an agent with zero money holdings is again denoted V_0. The value function of an agent with one unit of capital deposited at the bank is denoted V_1^d, and the value function of an agent carrying one unit of capital is denoted V_1^c. The corresponding value functions when the settlement system has suffered a disruption are given by $\tilde{V}_0$, $\tilde{V}_1^d$, and $\tilde{V}_1^c$, respectively. Finally, letting $V_1 = \max\{V_1^d, V_1^c\}$ and suppressing the arguments of ϕ, q, and τ, the value functions can be written:

$$\begin{aligned}
V_0 &= Kx\left\{-c+\beta\left[(1-q)V_1+q[\theta\tilde{V}_1^d+(1-\theta)\tilde{V}_1^c]\right]\right\} \\
&\quad +(1-Kx)\beta\left[(1-q)V_0+q\tilde{V}_0\right]+\tau \\
V_1^d &= (1-K)x\left\{u+\beta\left[(1-q)V_0+q\tilde{V}_0\right]\right\} \\
&\quad +[1-(1-K)x]\beta\left[(1-q)V_1+q\tilde{V}_1^d\right]+\tau+\phi \\
V_1^c &= (1-K)x\left\{u+\beta\left[(1-q)V_0+q\tilde{V}_0\right]\right\} \\
&\quad +[1-(1-K)x]\beta\left[(1-q)V_1+q\tilde{V}_1^c\right]+\tau
\end{aligned}$$

where

$$\begin{aligned}
\tilde{V}_0 &= (1-\theta)Kx(-c+\beta V_1)+[1-(1-\theta)Kx]\beta V_0 \\
\tilde{V}_1^d &= \beta V_1 \\
\tilde{V}_1^c &= (1-K)x(u+\beta V_0)+[1-(1-K)x]\beta V_1.
\end{aligned}$$

Consider V_0, the value function of a seller when the settlement system has not suffered a disruption. With probability K he meets a buyer, and with probability x the buyer likes his good and trade ensues: he incurs production cost c and receives payment. With probability θ this payment comes in the form of a check, and with probability $1-\theta$ he receives capital. If the payment system does not suffer a disruption at the end of the day, the seller becomes a buyer and has the choice of whether to store his capital at the bank or bring it with him to the market. That is, he begins the next day with value $V_1 = \max\{V_1^d, V_1^c\}$. However, if defaults are sufficiently large at the end of the current trading period, the settlement system shuts down and is not operational during the following day. If the seller had received a check, his funds are tied up in the banking system and cannot be withdrawn. In this case, he begins the following day with value $\tilde{V}_1^d$, which is simply the value of waiting a day to begin searching for a seller again. If, on the other hand, the seller received capital, he is able to trade the following period, and his value function is represented by $\tilde{V}_1^c$. A buyer with capital is able to trade, consume, and become a seller the following day (when the settlement system resumes operations). Finally, with probability $1-Kx$ the seller does not meet a buyer that likes his good, and thus he remains a seller. With probability $1-q$ there is no disruption, and the seller's value function is again V_0. On the other hand, with probability q a seller can only hope to trade with buyers holding capital, as buyers who utilized the banking sector cannot pay for current consumption. Finally, all agents receive a transfer, with expected value[20]

$$\tau(\mathcal{P}) = \bar{\sigma} B\kappa[\alpha(\mathcal{P};\mathcal{P})] - H(\mathcal{S}). \tag{20}$$

Similar reasoning can be employed to understand the remaining value functions.

An equilibrium in which agents use banks must again satisfy three incentive compatibility conditions: (9), (10), and $V_1^d \geq V_1^c$. Again, focusing on the case when $\theta = 1$, these conditions simplify to the following inequalities.[21]

$$\phi \geq \left[r + \frac{q(1-x)r}{1+r}\right] c - (1-K)x(u-c) \tag{21}$$

$$\phi \leq \left[r + \frac{q(1-x)r}{1+r}\right] u + Kx(u-c) \tag{22}$$

$$\phi \geq \frac{[1-(1-K)x](1-K)xq\left\{\left[r + x + \frac{q(1-x)r}{1+r}\right] u + Kxc\right\}}{(1+r)(r+x) + q\left[r(1-x) - [1-(1-K)x](1-K)x\right]}. \tag{23}$$

The first condition ensures that the interest paid on deposits is sufficiently large, and the probability of a disruption is sufficiently small, so that a seller is willing to incur the costs of production and accept a check. The second condition ensures that the interest earned on deposits is sufficiently small, or that there is a sufficiently high probability of a disruption, so that buyers have incentive to exchange their checks for goods, as opposed to leaving their deposits in the bank and simply consuming the interest. The final condition ensures that the interest on deposits is sufficiently large to compensate buyers for the risk of a disruption to the settlement system (and thus a period of autarky).

3.4 Equilibrium and Welfare

A banking equilibrium is again a pair (α^*, θ^*), with $\theta^* = 1$, such that: α^* satisfies the bank's optimality condition in (4) and uniquely determines ϕ^*, q^*, and τ^*; given these values the constraints (21) - (23) are satisfied; and $d = K/B$. Solving (14), given the closed-form solutions for V_0 and V_1 found in the appendix, we have

$$W = \frac{1+r}{r[1+q/(1+r)]}\left[\tau + M\phi + (1-K)xK(u-c)\right]. \tag{24}$$

[20]Note that the central bank only collects fees from those banks that have not failed.

[21]The derivation of these conditions, as well as closed-form solutions for the value functions, have been relegated to the appendix.

Plugging in (5), (19), and (20) yields

$$W(\mathcal{P}) \propto \frac{B\bar{\sigma}f\left\{[1-\alpha(\mathcal{P})]\,\frac{K}{B}\right\} - H(\mathcal{S}) + (1-K)Kx(u-c)}{1+[q(\mathcal{P})/(1+r)]}. \tag{25}$$

Note that the two policy levers $\mathcal{S}$ and k affect welfare in different ways. The cost of intra-day credit, k, only affects welfare through its influence on α, the level of reserves banks choose to retain; the actual amount of borrowing fees collected is simply a transfer and has no additional consequences for welfare. The frequency of settlement, on the other hand, has substantive effects on welfare above and beyond its effect on α. In particular, for a fixed α, increasing the frequency of settlement implies an increase in the resource costs and a decrease in potential defaults via the accumulation effect.

For each settlement frequency $\mathcal{S}$, then, there exists at least one reserve ratio $\alpha_{\mathcal{S}} \in [0,1]$ that maximizes welfare. Therefore, choosing the optimal settlement system is tantamount to choosing the settlement frequency $\mathcal{S}$ that maximizes welfare given k is chosen to achieve the optimal level of reserves $\alpha_{\mathcal{S}}$. In the next section, we study this question of optimal settlement system design within the context of a numerical example.

4 An Illustrative Example

In the general framework presented above, we established a trade-off between alternative settlement system policies and characterized the welfare function. In this section, we offer an example to illustrate several insights that the model provides about how, precisely, the optimal settlement system policy depends on specific features of the economic environment, such as resource costs, the fragility of the settlement system, the likelihood of bank failure, the volume of trade, and the productivity of capital investments.

4.1 A Simple Environment

Suppose that there are only two sub-periods in the day, so that $T = 2$, and consider the following two settlement systems. The first, $\mathcal{S}' = \{2\}$, is a *deferred* settlement system, as transactions occurring at $t = 1$ are not finalized until the end of the day (at $t = 2$). The second settlement system, $\mathcal{S}'' = \{1, 2\}$, is a *real-time* settlement system, as transactions are finalized as they take place.

Let the production technology be given by $f(i) = z \cdot i^a$, with $z > 0$ and $a < 1$, and let σ be drawn from the cumulative density function $G(\sigma) = \sigma^n$ for some $n > 1$. Note that this implies an expected value $\bar{\sigma} = n/(n+1)$.

The matching technology in the trading sector implies that the realization of net liabilities, l, is distributed according to the following process.[22] At $t = 1$, three events are possible. First, with probability ρ, half of a bank's buyers make purchases and no sellers make any sales, so that $l_1 = d/2$. With equal probability ρ, half of a bank's sellers make sales and buyers are inactive, so that $l_1 = -d/2$. Finally, with probability $1-2\rho$, a bank's customers do not engage in trade, and $l_1 = 0$. A similar trend occurs in the afternoon: either a bank's buyers are active and its sellers are inactive, its sellers are active and its buyers are inactive, or both are inactive. In particular, if $l_1 = d/2$, then at $t = 2$ net liabilities return to zero with probability $1 - \rho$ and increase to d with probability ρ. Similarly, if $l_1 = -d/2$, then at $t = 2$ net liabilities return to zero with probability $1 - \rho$ and decrease to $-d$ with probability ρ. Finally, if $l_1 = 0$, then at $t = 2$ net liabilities remain at zero with probability $1 - 2\rho$, and increase or decrease to d with probability ρ. Figure 1 below illustrates the distribution of net liabilities.

This formulation implies a relationship between economic activity in the trading sector and the evolution of net liabilities for a representative bank. In particular, given that $d = K/B$ in any equilibrium with $\theta = 1$, we have

$$B\left\{2\left[\rho(1-\rho)+\rho^2+(1-2\rho)\rho\right]\frac{K}{B}\right\} = 2K(1-K)x. \tag{26}$$

Equation (26) reflects that the aggregate number of payments processed in the banking sector must equal the number of purchases and sales executed in the trading sector. For any value of K, equation (26) determines the relationship between the volatility of intra-day liabilities in the banking sector (ρ) and the amount of economic activity in the trading sector (x).

Given $\lambda(l)$, the resource costs are easily characterized:

$$H(\mathcal{S}') \equiv H_D = 2\left[\rho^2+(1-2\rho)\rho\right]hB = 2\rho(1-\rho)hB \tag{27}$$

$$H(\mathcal{S}'') \equiv H_R = 2\left[2\rho^2+2\rho(1-\rho)+(1-2\rho)\rho\right]hB = 2\rho(3-2\rho)hB \tag{28}$$

where the subscripts "D" and "R" denote the deferred and real-time settlement frequencies, respectively. Equation (27) illustrates that the only banks that require the central bank to finalize payment are those with realizations of l

[22] Recall that we have assumed that agents are matched once each day, and that we have allowed for the timing of that match and the banks of each agent to be potentially correlated. Since buyers and sellers are not at all affected by the sub-period at which they meet, or the bank of their trading partner, it is not necessary to write down a precise matching process. However, note that it is not difficult to write down a specification that would generate the distribution $\lambda(\cdot)$ described here.

Figure 1: Distribution of Net Liabilities

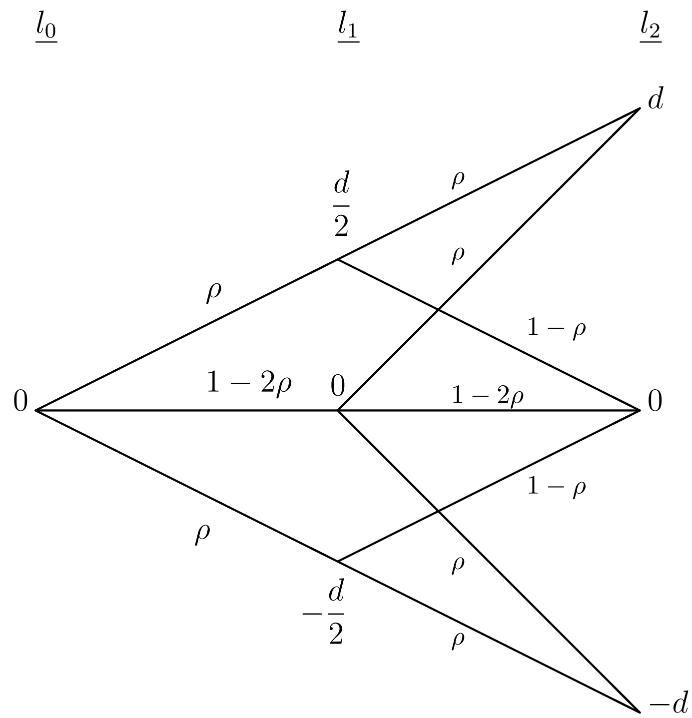

such that $l_2 = d$ or $l_2 = -d$. On the other hand, in a real-time settlement system the central bank must transfer payments once to or from a bank that realizes $l = (0, -d)$ or $l = (0, d)$, respectively, and twice for the realizations $l = (\frac{d}{2}, d)$, $l = (\frac{-d}{2}, -d)$, $l = (\frac{d}{2}, 0)$ and $l = (\frac{-d}{2}, 0)$.

Finally, let $\bar{\delta} = \chi K$, so that χ represents the minimum fraction of total capital owed by failed banks that would cause a disruption to the settlement system.

4.2 Welfare and Optimal Settlement System Policy

We now study the equilibrium level of welfare given the parametric forms specified above. We restrict our analysis to parameter values for which the optimal reserve ratio under both settlement systems is less than $\frac{1}{2}$.[23] Note that for $\alpha \in [0, 1/2]$,

$$\delta_D(\sigma) = B(1-\sigma)\left[\rho(1-\rho)(1-\alpha)d\right] \quad (29)$$

$$\delta_R(\sigma) = B(1-\sigma)\left[(1-2\rho)\rho(1-\alpha)d + \rho^2(d/2)\right]. \quad (30)$$

From (29), we see that aggregate defaults are equivalent to the outstanding net liabilities of failed banks that realized either $l = (0, d)$ or $l = (\frac{d}{2}, d)$. For either realization, the value of outstanding liabilities in excess of reserves is $d - \alpha d$. From (30), note that under a real-time system, settling at $t = 1$ limits the accumulation of net liabilities for the realization $l = (\frac{d}{2}, d)$ (the accumulation effect). The expected value of aggregate default is simply $\Delta_j = \delta_j(\bar{\sigma})$ for $j \in \{D, R\}$. Moreover, given $\bar{\delta} = \chi K$, we can solve for $\underline{\sigma}_j$ to find

$$q_D = \left[1 - \frac{\chi K}{B\left[\rho(1-\rho)(1-\alpha)d\right]}\right]^n \quad (31)$$

$$q_R = \left[1 - \frac{\chi K}{B\left[(1-2\rho)\rho(1-\alpha)d + \rho^2(d/2)\right]}\right]^n. \quad (32)$$

It is easy to show, as proven more generally in Section 2.1, that $\Delta_D \geq \Delta_R$ and $q_D \geq q_R$ for any α. From (15), then, we can plug in H_j from (27) and (28), $\delta_j(\bar{\sigma})$ from (29) and (30), and q_j from (31) and (32) to get a well-defined welfare function $W_j(\alpha)$ for $j \in \{D, R\}$. Let $\alpha_j^* \in \arg\max_{\alpha \in [0,1/2]} W_j$.[24] Using (4), we can characterize the bank's profit-maximizing reserve ratio as a function of the borrowing cost:

$$\alpha_D(k) = 1 - \left[\frac{\rho(1-\rho)k}{azd^{a-1}}\right]^{\frac{1}{a-1}} \quad (33)$$

$$\alpha_R(k) = 1 - \left[\frac{2\rho(1-\rho)k}{azd^{a-1}}\right]^{\frac{1}{a-1}}. \quad (34)$$

Therefore, the optimal intra-day credit policy for settlement frequency $j \in \{D, R\}$ is simply the k_j^* such that $\alpha_j(k_j^*) = \alpha_j^*$. Though closed-form characterizations of α_j^* are not feasible, we proceed numerically to characterize the optimal intra-day credit policies and the welfare-maximizing joint policy (a deferred system with k_D^* or a real-time system with k_R^*) conditional

[23]If $\alpha_D \in [\frac{1}{2}, 1]$, then $\alpha_D = \alpha_R$, and we lose two margins of interest: interest rates are identical under both settlement frequencies and the reserve effect vanishes.

[24]Since W_j is continuous on α, we are assured of the existence of such an α_j^*.

on the fundamentals of the economy. In particular, we specify a set of benchmark parameter values, solve for the optimal settlement system numerically, and then perturb those values in order to characterize features of an economy that make one system preferred to the other. Note that the values chosen are not a calibration to any particular economy per se, but instead are simply meant to represent reasonable dynamics.

To that end, let $\rho = .3$, so that there is some intra-day volatility of net liabilities, and also some mean-reversion in l. Given $M = .5$, this implies $x = .84$. Moreover, let $n = 499$, so that the mean fraction of failed banks in any given period is 0.2%. We let $z = .5$, $a = .5$, and $h = .2$. These values imply, at the optimal borrowing policy, that resource costs are approximately 1% - 2% of the total value of revenue generated from banks' investments. We let $B = .005$, so that each bank begins with one hundred depositors, and $\xi = .00175$, which implies a settlement system disruption approximately every 200 periods. Finally, we set $r = .01$, $u = 1$, and $c = .1$.[25]

Under these parameter values, we can solve for the optimal level of reserves for each settlement frequency, α_j^*, and the corresponding intra-day credit policy k_j^* that the central bank can set to implement this portfolio allocation. For example, in figures 2 and 3 below, we plot α_j^* and k_j^* across different values of z, which represents the productivity of the investment technology available to banks; the solid line corresponds to $j = D$ and the dashed line corresponds to $j = R$. There are three things to notice. First of all, we find that the optimal intra-day credit policy is *not* necessarily to provide free liquidity. Since banks do not internalize the negative externality associated with systemic risk (i.e. a disruption the settlement system), there are welfare gains from the central bank using the price of liquidity to promote more conservative reserve management. Secondly, note that the optimal intra-day credit policy is different across settlement frequencies. This implies, again, that the question of optimal settlement design is a *joint* question about settlement frequency and credit policy. Finally, we find here (and for most parameter values) that $k_D^* \geq k_R^*$. This is natural, as banks have greater liquidity needs under a real-time settlement frequency, and hence reserve management is more sensitive to the cost of credit.

[25] Setting u much larger than c insures the incentive constraints are met. In all equilibrium computed in this section, the inequalities (21) - (23) are satisfied.

Figure 2: Optimal Reserve Ratio (α^*)

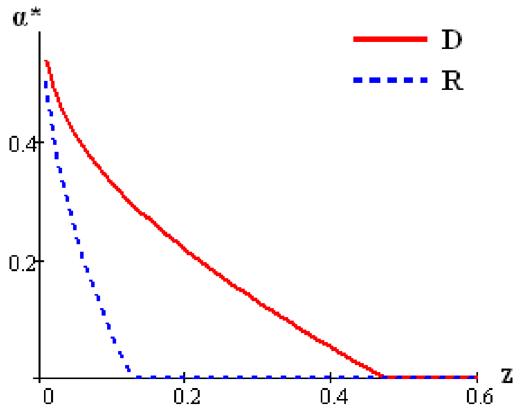

Figure 3: Optimal Credit Policy (k^*)

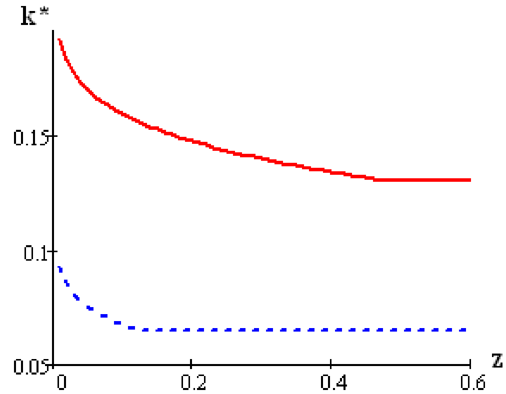

Moreover, increased settlement frequency reduces risk independently of the reserve ratio via the accumulation effect, thereby reducing the need for the central bank to curb risk through the intra-day credit rate. This can be seen in figure 4: even when $\alpha_D^* = \alpha_R^* = 0$, $q_D^* > q_R^*$ because, under settlement system R, settlement at $t = 1$ has limited the total amount of potential outstanding liabilities.

Figure 4: Probability of Disruption (q^*)

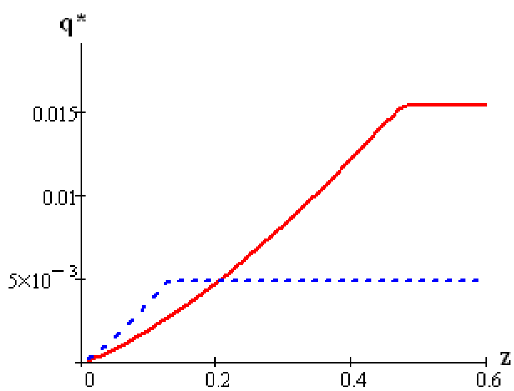

Given the optimal intra-day lending policy k_j^* for $j \in \{D, R\}$, we can now determine the optimal settlement system $(\mathcal{S}, k)$ conditional on the fundamentals of the economy. As noted earlier, over the past two decades many countries have changed from a deferred to real-time settlement system. It was argued that the costs of a real-time system had become sufficiently low, while the risks of deferred settlement had grown as a result of financial integration and increased exposure to external shocks. The current model lends at least qualitative support to this argument.

Figures 5 and 6 plot the regions of the parameter space in which maximal welfare is achieved by either the optimal real-time (R) or deferred (D) settlement system. In figure 5, we see that a real-time settlement system with the optimal corresponding intra-day credit policy is preferred when systemic risk is high (χ is low) or when the likelihood of bank failure ($1-\bar{\sigma}$) is relatively large. Figure 6 provides two additional insights. First, and perhaps not surprisingly, the optimal deferred settlement system is preferred when resource costs (h) are large, since infrequent settlement decreases the number of payments that the central bank must finalize.

Figure 5: Optimal System: χ vs. $1-\bar{\sigma}$

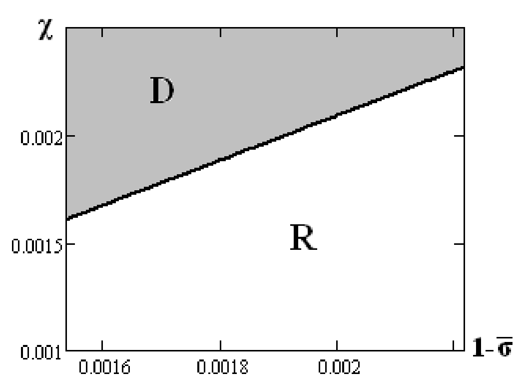

Figure 6: Optimal System: h vs. x

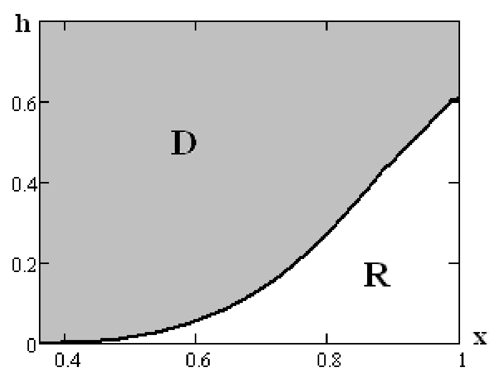

Less obvious is the relationship between the optimal settlement system and the volume of trade in the trading sector, given here by x, which uniquely determines ρ through (26). It should be noted that this relationship would be very difficult to study in existing models, which typically do not model this link between activity in the banking sector and activity in the real sector. What we find is that when x is relatively small, the volatility of intra-day liabilities is small, and hence the risk of default is low. In this case, a deferred settlement system is preferred. Alternatively, in economies with a large volume of trade and more volatile intra-day liabilities, a real-time system is typically optimal.

5 Conclusion

Efficient, safe settlement systems are vital to modern economies. As the value of transfers processed has increased dramatically in recent years, thus magnifying the risks of a settlement system failure, policymakers throughout the world have chosen to re-evaluate their procedures. In particular, settlement system operators have sought to find the optimal settlement system frequency and intra-day credit policy to balance costs and risks.

However, despite the importance of settlement systems, the topic remains under-researched in the academic literature and many critical questions remain unanswered. The current paper attempts to formalize some of the issues at hand, and make progress on certain specific questions. First, we have formalized the widely discussed trade-off between cost and risk associated with settlement system design. Secondly, we have integrated this model of banks into an economy with real production, trade, and consumption. Using this model, we are able to characterize the optimal settlement system design as a function of the economy's fundamentals. Some results are straight-forward: an optimal system with greater settlement frequency is preferred when resource costs are low and risks are high. However, the model also reveals a more sophisticated relationship between optimality and the level of economic activity in the real sector; we find that an optimal system with greater settlement system is preferred when the level of economic activity is high.

The basic model developed here can easily be extended to address other relevant questions concerning settlement system design. A current area of interest is the effect of interest payments on reserves held at the central bank. Moreover, the model could be extended to include inter-day dynamics (in addition to the currently modeled intra-day dynamics) in order to assess the optimality of overnight borrowing policies. Lastly, by modeling the informational asymmetries between the central bank and financial institutions, one could make formal progress on the optimal bail-out policies that the central bank should institute as the "lender of last resort". These are areas for future work.

Appendix

We first establish a series of intermediate results.

Corollary 1. *For any l, α, $t \in \mathcal{T}$, and $\mathcal{S} = \{t_1, \ldots, t_F\}$,*

$$\hat{l}_t(l, \alpha; \mathcal{S}) = \max\Big\{0, \min\{l_t - l_{s_t}, l_t - l_{s_{t-1}}, \ldots, l_t - l_{t_1}, l_t - \alpha d\}\Big\}, \tag{35}$$

where s_t corresponds to the period in which the most recent settlement occurred prior to t, s_{t-1} the settlement prior to s_t, and so on.

The proof follows immediately from recursive substitution.

Corollary 2. *For any l, α, and $\mathcal{S} = \{t_1, \ldots, t_F\}$,*

$$b(l, \alpha; \mathcal{S}) = \max\{0, l_{t_1} - \alpha d, l_{t_2} - \alpha d, \ldots, l_{t_F} - \alpha d, \}. \tag{36}$$

Proof. We sketch the proof, which proceeds by induction. Recall that

$$b(l, \alpha; \mathcal{S}) = \hat{l}_{t_1} + \hat{l}_{t_2} + \cdots + \hat{l}_{t_F},$$

so that it is sufficient to show, for any finite $N \in \mathbb{N}$,

$$\hat{l}_{t_1} + \hat{l}_{t_2} + \cdots + \hat{l}_{t_N} = \max\{0, l_{t_1} - \alpha d, l_{t_2} - \alpha d, \ldots, l_{t_N} - \alpha d, \}. \tag{37}$$

The first step of the proof by induction is trivial, as (37) is true by definition. Now we assume that this is true for an arbitrary N, and show that it is then true for $N+1$. If

$$\hat{l}_{t_1} + \hat{l}_{t_2} + \cdots + \hat{l}_{t_N} = \max\{0, l_{t_1} - \alpha d, l_{t_2} - \alpha d, \ldots, l_{t_N} - \alpha d, \}$$

then using Corollary 1

$$\hat{l}_{t_1} + \cdots + \hat{l}_{t_N} + \hat{l}_{t_{N+1}} = \max\{0, l_{t_1} - \alpha d, \ldots, l_{t_N} - \alpha d\} + \max\Big\{0, \min\{l_{t_{N+1}} - l_{t_N}, l_{t_{N+1}} - l_{t_{N-1}}, \ldots, l_{t_{N+1}} - l_{t_1}, l_{t_{N+1}} - \alpha d\}\Big\}.$$

Suppose $\exists\, t_i \in \{t_1, \ldots, t_N\}$ such that $l_{t_i} \geq l_{t_{N+1}}$, then

$$\max\Big\{0, \min\{l_{t_{N+1}} - l_{t_N}, l_{t_{N+1}} - l_{t_{N-1}}, \ldots, l_{t_{N+1}} - l_{t_1}, l_{t_{N+1}} - \alpha d\}\Big\} = 0$$

and

$$\begin{aligned}\hat{l}_{t_1} + \cdots + \hat{l}_{t_N} + \hat{l}_{t_{N+1}} &= \max\{0, l_{t_1} - \alpha d, \ldots, l_{t_N} - \alpha d\} \\ &= \max\{0, l_{t_1} - \alpha d, \ldots, l_{t_N} - \alpha d, l_{t_{N+1}} - \alpha d\}.\end{aligned}$$

Conversely, suppose $l_{t_{N+1}} > l_{t_i}\ \forall t_i \in \{t_1,\ldots,t_N\}$ and let $l_{t^*} = \max\{l_{t_1},\ldots,l_{t_N}\}$. Then

$$\begin{aligned}\hat{l}_{t_1} + \cdots + \hat{l}_{t_N} + \hat{l}_{t_{N+1}} &= \max\{0, l_{t^*} - \alpha d\} + \max\Big\{0, \min\{l_{t_{N+1}} - l_{t^*}, l_{t_{N+1}} - \alpha d\}\Big\} \\ &= \max\{0, l_{t_{N+1}} - \alpha d\} \\ &= \max\{0, l_{t_1} - \alpha d, \ldots, l_{t_N} - \alpha d, l_{t_{N+1}} - \alpha d\}.\end{aligned}$$

■

Corollary 3. *At any settlement frequency, expected borrowing fees exhibit convexity in* α*. That is,* $\forall\ \alpha_1, \alpha_2 \in [0,1],\ \alpha_1 > \alpha_2\ \Rightarrow\ \kappa'(\alpha_1;\mathcal{P}) \geq \kappa'(\alpha_2;\mathcal{P})$.

Proof. From Corollary 2, it is clear that the function $\alpha \mapsto b(l,\alpha;\mathcal{S})$ is the upper envelope of a collection of affine functions of α, which is well-known to be convex.[26] Since b is convex in α, then clearly κ is as well. ■

Corollary 4. *The impact of a marginal increase in the reserve ratio is larger (more negative) under a higher frequency of settlement. That is, for any* $\alpha \in [0,1]$, $\kappa'(\alpha;\mathcal{P}'_{\mathcal{S}}) \geq \kappa'(\alpha;\mathcal{P}''_{\mathcal{S}})$.

Proof. Ignoring the points of discontinuity, it follows from Corollary 2 that

$$0 \geq \frac{\partial b(l,\alpha;\mathcal{S}')}{\partial \alpha} \geq \frac{\partial b(l,\alpha;\mathcal{S}'')}{\partial \alpha}.$$

Again, taking expectations yields the desired result. ■

Proof of Lemma 2. From Corollary 2, it is obvious that $b(l,\alpha;\mathcal{S})$ is weakly decreasing in α for all l, and strictly decreasing in α for l such that $b(l,\alpha;\mathcal{S}) > 0$. Moreover, if $\alpha d < \bar{L}$, then $b(l,\alpha;\mathcal{S}) > 0$ with strictly positive probability. Therefore, it follows that $\kappa(\alpha;\mathcal{P}) = \mathbb{E}_l\Big[k \cdot b(l,\alpha;\mathcal{S})\Big]$ is weakly decreasing in α, and strictly decreasing if $\alpha < \frac{\bar{L}}{d}$. ■

Proof of Lemma 3. From Corollary 2, it is clear that $b(l,\alpha;\mathcal{S}') \leq b(l,\alpha;\mathcal{S}'')$ for any l and α, so that $\kappa(\alpha;\mathcal{P}'_{\mathcal{S}}) \leq \kappa(\alpha;\mathcal{P}''_{\mathcal{S}})$. We will now show that $\alpha(\mathcal{P}'_{\mathcal{S}}) \leq \alpha(\mathcal{P}''_{\mathcal{S}})$.

Suppose, towards a contradiction, that $\alpha(\mathcal{P}'_{\mathcal{S}}) > \alpha(\mathcal{P}''_{\mathcal{S}})$. By the strict concavity of f, this implies that

$$f'\left((1-\alpha(\mathcal{P}'_{\mathcal{S}}))d\right) > f'\left((1-\alpha(\mathcal{P}''_{\mathcal{S}}))d\right). \tag{38}$$

[26] I thank an anonymous referee for suggesting this simple proof.

By the bank's first order condition (4), this implies

$$\kappa'\left(\alpha(\mathcal{P}''_{\mathcal{S}});\mathcal{P}''_{\mathcal{S}}\right) > \kappa'\left(\alpha(\mathcal{P}'_{\mathcal{S}});\mathcal{P}'_{\mathcal{S}}\right). \tag{39}$$

But from Corollary 4 we have

$$\kappa'\left(\alpha(\mathcal{P}''_{\mathcal{S}});\mathcal{P}'_{\mathcal{S}}\right) \geq \kappa'\left(\alpha(\mathcal{P}''_{\mathcal{S}});\mathcal{P}''_{\mathcal{S}}\right), \tag{40}$$

and from Corollary 3 we have

$$\kappa'\left(\alpha(\mathcal{P}'_{\mathcal{S}});\mathcal{P}'_{\mathcal{S}}\right) \geq \kappa'\left(\alpha(\mathcal{P}''_{\mathcal{S}});\mathcal{P}'_{\mathcal{S}}\right), \tag{41}$$

so that

$$\kappa'\left(\alpha(\mathcal{P}'_{\mathcal{S}});\mathcal{P}'_{\mathcal{S}}\right) \geq \kappa'\left(\alpha(\mathcal{P}''_{\mathcal{S}});\mathcal{P}''_{\mathcal{S}}\right), \tag{42}$$

a contradiction. The proof that $\kappa(\alpha;\mathcal{P}'_k) \leq \kappa(\alpha;\mathcal{P}''_k)$ and thus $\alpha(\mathcal{P}'_k) \leq \alpha(\mathcal{P}''_k)$ follows in an identical fashion, given that $\kappa'(\alpha;\mathcal{P}'_k) \geq \kappa'(\alpha;\mathcal{P}''_k)$. ■

Proof of Proposition 1. Given $\alpha(\mathcal{P}'_j) \leq \alpha(\mathcal{P}''_j)$, it is trivial to show that $Y(\mathcal{P}'_j) \geq Y(\mathcal{P}''_j)$. Suppose, towards a contradiction, that $\phi(\mathcal{P}'_j) < \phi(\mathcal{P}''_j)$. Since $\alpha(\mathcal{P}'_j)$ is optimal, it follows that

$$\phi(\mathcal{P}'_j) = \frac{1}{d}\Big[\bar{\sigma}f[(1-\alpha(\mathcal{P}'_j))d]-\kappa(\alpha(\mathcal{P}'_j);\mathcal{P}'_j)\Big] \geq \frac{1}{d}\Big[\bar{\sigma}f[(1-\alpha(\mathcal{P}''_j))d]-\kappa(\alpha(\mathcal{P}''_j);\mathcal{P}'_j)\Big].$$

Then if

$$\phi(\mathcal{P}''_j) = \frac{1}{d}\Big[\bar{\sigma}f[(1-\alpha(\mathcal{P}''_j))d]-\kappa(\alpha(\mathcal{P}''_j);\mathcal{P}''_j)\Big] > \frac{1}{d}\Big[\bar{\sigma}f[(1-\alpha(\mathcal{P}'_j))d]-\kappa(\alpha(\mathcal{P}'_j);\mathcal{P}'_j)\Big] = \phi(\mathcal{P}'_j),$$

we have that

$$\kappa(\alpha;\mathcal{P}'_j) > \kappa(\alpha;\mathcal{P}''_j),$$

a contradiction to Lemma 3. ■

Proof of Lemmas 4 and 5. These results follow immediately from Corollary 1. ■

Derivation of Value Functions and Incentive Constraints. For notational simplicity, let $\xi \equiv 1 - 1/(1+r)$. One can solve the six value functions simultaneously, imposing the equilibrium conditions that $\theta = 1$ and $V_1 = V_1^d$, to find that:

$$\begin{aligned} V_0 &= \frac{1+r}{r+q\xi}\left\{\tau + \frac{Kx(1-q\xi)[\phi+(1-K)xu] - [r+(1-K)x+q\xi[1-(1-K)x]]Kxc}{r+x+(1-x)q\xi}\right\} \\ V_1 &= \frac{1+r}{r+q\xi}\left\{\tau + \frac{[r+Kx+q\xi(1-Kx)][\phi+(1-K)xu] - (1-K)x(1-q\xi)Kxc}{r+x+(1-x)q\xi}\right\}. \end{aligned}$$

Given these two expressions, one can calculate

$$V_1 - V_0 = \frac{(1+r)[\phi + (1-K)Xu + Kxc]}{r + x + (1-x)q\xi}.$$

From this expression, one can easily derive conditions on ϕ so that (9) and (10) hold. Lastly, subtracting V_1^d - V_1^c yields

$$V_1^d \geq V_1^c \;\Leftrightarrow\; \phi \geq [1-(1-K)x]q\beta(1-K)x[u + \beta(V_0 - V_1)],$$

which allows for the derivation of conditions under which $V_1^d \geq V_1^c$. ■

References

M.L. Bech and B. Hobijn. Technology Diffusion within Central Banking: The Case of Real-Time Gross Settlement. *International Journal of Central Banking*, 3(3):147–182, 2007.

A.N. Berger, D. Hancock, and J.C. Marquardt. A Framework for Analyzing Efficiency, Risks, Costs, and Innovations in the Payments System. *Journal of Money, Credit and Banking*, 28(4):696–732, 1996.

B.S. Bernanke. Clearing and settlement during the crash. *Review of Financial Studies*, 3(1):133–151, 1990.

R.O. Cavalcanti and N. Wallace. Inside and outside Money as Alternative Media of Exchange. *Journal of Money, Credit & Banking*, 31(3), 1999a.

R.O. Cavalcanti and N. Wallace. A Model of Private Bank-Note Issue. *Review of Economic Dynamics*, 2(1):104–136, 1999b.

M.J. Fleming and K.D. Garbade. When the Back Office Moved to the Front Burner: Settlement Fails in the Treasury Market after 9/11. *Federal Reserve Bank of New York Economic Policy Review*, 8(2):35–57, 2002.

S. Freeman. The Payments System, Liquidity, and Rediscounting. *The American Economic Review*, 86(5):1126–1138, 1996.

S. Freeman. Rediscounting under aggregate risk. *Journal of Monetary Economics*, 43(1):197–216, 1999.

P. He, L. Huang, and R. Wright. Money and Banking in Search Equilibrium. *International Economic Review*, 46(2):637–670, 2005.

D.B. Humphrey, L.B. Pulley, and J.M. Vesala. Cash, Paper, and Electronic Payments: A Cross-Country Analysis. *Journal of Money, Credit and Banking*, 28(4):914–39, 1996.

C. Kahn and W. Roberds. Payment System Settlement and Bank Incentives. *Review of Financial Studies*, 11(4):845–870, 1998.

C.M. Kahn, J. Mcandrews, and W. Roberds. Settlement Risk under Gross and Net Settlement. *Journal of Money, Credit & Banking*, 35(4):591–609, 2003.

N. Kiyotaki and R. Wright. A search-theoretic approach to monetary economics. *The American Economic Review*, pages 63–77, 1993.

N.R. Kocherlakota. Money Is Memory. *Journal of Economic Theory*, 81(2): 232–251, 1998.

T. Koeppl, C. Monnet, and T. Temzelides. Mechanism Design and Payments. Technical report, Mimeo, University of Pittsburgh, 2004.

T. Koeppl, C. Monnet, and T. Temzelides. A Dynamic Model of Settlement. *Journal of Economic Theory*, 142(1):233–246, 2008.

A. Martin. Optimal pricing of intraday liquidity. *Journal of Monetary Economics*, 51(2):401–424, 2004.

A. Martin. Recent Evolution of Large-Value Payment Systems: Balancing Liquidity and Risk. *Economic Review, Federal Reserve Bank of Kansas City, First Quarter*, 2005.

A. Martin and J. McAndrews. Liquidity-saving mechanisms. *Journal of Monetary Economics*, 55(3):554–567, 2008.

J. McAndrews and S. Potter. Liquidity Effects of the Events of September 11, 2001. *Economic Policy Review, Federal Reserve Bank of New York*, 8(2), 2002.

D.C. Mills. Alternative Central Bank Credit Policies for Liquidity Provision in a Model of Payments. *Journal of Monetary Economics*, 53(7):1593–1611, 2006.

G. Selgin. Wholesale payments: questioning the market-failure hypothesis. *International Review of Law & Economics*, 24(3):333–350, 2004.

N. Wallace. Whither Monetary Economics? *International Economic Review*, 42(4):847–869, 2001.

R. Zhou. Understanding Intraday Credit in Large-Value Payment Systems. *Federal Reserve Bank of Chicago Economic Perspectives*, 24(3):29–44, 2000.