Efficiency and Fragility in Loss Networks

Johan H. Ugander
University of Cambridge
Part III Essay

May 2009

Contents

Introduction

In this essay, we study the stochastic dynamics of resource management and lending from a network traffic perspective. We find that as the economic efficiency of lending management strategies is increased, the network agents become increasingly exposed to systemic risks.

We first review the standard model of independent resource capacity management due to Erlang, initially developed for telecommunications systems. In the second section, we motivate the applicability of this resource modeling approach to liquidity. These first two sections serve as background for the third section, where we develop a modeling framework for resource lending networks utilizing the theory of loss networks and alternative routing, presenting results for symmetric lending networks whereby operational costs are greatly decreased through resource lending.

In the fourth section, we present our main results, studying the systemic connectivity brought on by lending and the effect of lending defaults, which breaks the symmetric benefit of lending participation. As default risk is increased, we observe a rapid transition in optimal lending strategy similar to phase transitions in physical systems, where the favorability of lending participation transitions discontinuously from good to bad at a critical probability of default. Using these results, we propose how a small change in default risk can lead to a systemic liquidity crisis in the greater process of bank lending.

In a concluding section, we propose directions for future research, primarily towards the study of heterogenous lending networks. We comment briefly on the possible applications of game theory to lending participation.

1 Independent resource management

In the classic telecommunications setting [1], resource management refers to telephone link utilization by customer calls arriving as a Poisson process, initiated at a constant rate that is uniform in time, and utilizing links for an exponentially distributed period of time.

1.1 The birth and death process

The utilization of C resource units can then be modeled by a finite birth and death process, which is truncated from above by the capacity C. The transition rates become

$$\begin{array}{rcll} q(j,j+1) & = & \lambda, & j=0,1,\ldots,C-1 \\ q(j,j-1) & = & j\mu, & j=1,2,\ldots,C \end{array}, \tag{1}$$

where λ is the rate of the arrival process and μ is the rate of the departure process. Through detailed balance and normalization, it can be shown that the probability distribution $\pi_C(j)$ (subscript C to emphasize that different capacities lead to different equilibrium distributions) becomes

$$\pi_C(j) \quad = \quad \frac{\nu^j}{j!}\left(\sum_{i=0}^{C}\frac{\nu^i}{i!}\right)^{-1}, \qquad j=0,1,\ldots,C, \tag{2}$$

where $\nu = \lambda/\mu$. The probability of a call being blocked (the *blocking probability*) then takes the form of the classic Erlang-B formula,

$$E(\nu, C) = \pi_C(C) = \frac{\nu^C}{C!} \left(\sum_{i=0}^{C} \frac{\nu^i}{i!} \right)^{-1}. \tag{3}$$

1.2 Cost minimization

The operational challenge for the resource manager is to determine the optimal capacity C where profits are maximized. We assume that there exists a contractual obligation to meet all withdrawal requests, and that unsatisfied requests must be satisfied through more expensive external means.

When managing a pool of equipment, this may mean temporarily leasing equipment directly from the manufacturer. To provide an example, many services models for 'utility computing', where mainframe servers are sold and leased to computing customers, offer to sell computing systems with extra idle processors at no significantly elevated purchasing price. These extra processors are leased temporarily 'on-demand' at a later date, to manage peak loads, albeit for a substantial fee. The question for the computing customer is then to decide how many processors to purchase directly and how many to leave for emergency leasing.

Without resource lending, this operational challenge is then a matter of balancing the opportunity cost of owning C resource units with the costs of emergency leasing (which occurs for the fraction of requests that are blocked, $E(\nu, C)$). The optimization problem becomes

$$\begin{aligned} C_{opt}(\nu) &= \underset{C \in \mathbb{Z}^+}{\operatorname{argmin}} \; \omega_{INV} C + \omega_{EXT} \nu E(\nu, C), & (4) \\ &= \underset{C \in \mathbb{Z}^+}{\operatorname{argmin}} \; \underbrace{C + \frac{\omega_{EXT}}{\omega_{INV}} \nu E(\nu, C)}_{\mathcal{C}_1(C)}, & (5) \end{aligned}$$

where ω_{INV} is the opportunity cost interest rate, the alternative return on investment for the value of a resource unit elsewhere, and ω_{EXT} is the interest rate for emergency leasing.

An example of the cost function $\mathcal{C}_1(C)$ for a fixed traffic rate ($\nu = 100$) under three diffferent interest rate scenarios is shown in Figure 1(a), while the optimal capacity C_{opt} as a function of ν is shown in Figure 1(b), for the same three scenarios.

(a)

(b)

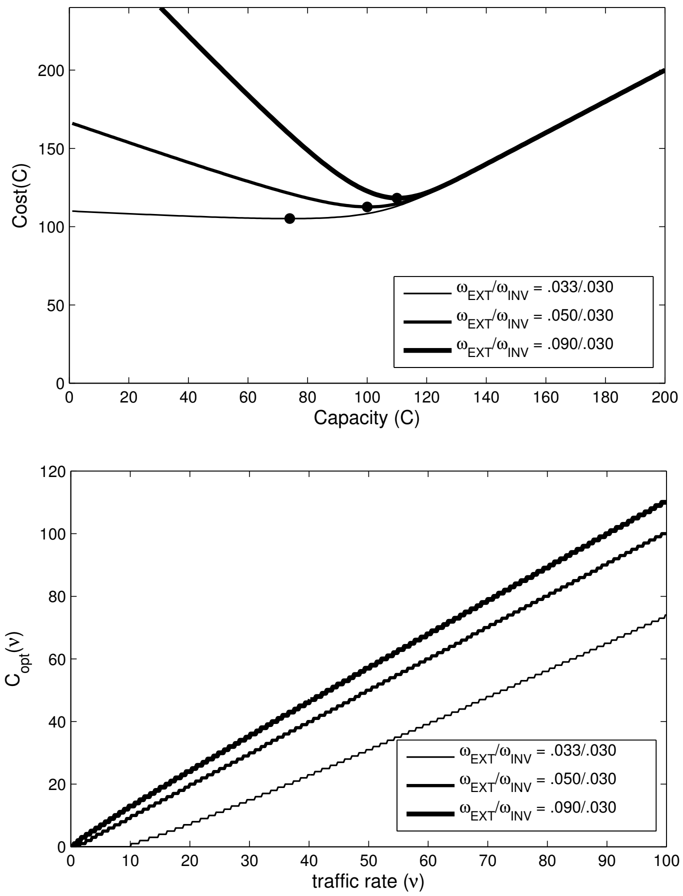

Figure 1: (a) The cost function $\mathcal{C}_1(C)$ from (5), calculated for three different interest rate scenarios, with the differnet optima (C_{opt}) indicated. The scenario $\omega_{EXT}/\omega_{INV} \leq 1$ would result in $C_{opt} = 0$, since requests might as well be shifted to the external mechanism. Notice that under certain interest rate conditions, it may be optimal to hold fewer resource units than are being utilized on average, owing to the cheap availability of external resources. (b) The optimal capacities $C_{opt}(\nu)$, for the same three interest rate scenarios. Capacities are restricted to whole units.

2 Liquidity as a resource

The purpose of this section is to accommodate the specific details of liquidity management within the context of the vanilla resource management problem as presented in the previous section. By considering a more technically correct model and offering a comparison, we demonstrate that optimal liquidity management can be well approximated by optimal vanilla resource management.

As a vanilla resource management problem, we may view bank lending as a process whereby banks hold C units of currency reserved for liquidity management, and these units are withdrawn through a withdrawal process with events uniformly distributed in time, and returned after an exponentially distributed period of time. In this case we have precisely the process described in (1), where state j refers to having j of the C units currently withdrawn. This is, however, a mildly inaccurate view of process.

More accurately, the process is the 'reverse' of (1), where it is the deposit arrivals that are uniformly distributed in time, and the withdrawal process intensity is linear in the amount of liabilities held with the bank. We now aim to show that the two models are truncations of opposite tails of the same nearly symmetric distribution, and therefore exhibit identical dynamics to a very close approximation.

2.1 Lower truncated liquidity process

Let us assume that a bank receives stationary streams of deposit and withdrawal traffic. Such a bank will then aim to invest its liabilities for a higher return than promised to the depositor, through loans or other asset purchases. Banks can not invest all of their funds, because they must maintain a certain level of liquid assets in order to meet withdrawal traffic. This avoids being forced to sell assets, a central hazard in banking because often times such assets are invested with contractual obligations, and can not be made available to customers without a financial penalty to the bank for premature termination. Furthermore, potentially illiquid assets may be difficult to sell quickly for a fair price, since they may not be so regularly traded.

And so, banks maintain a certain amount of liquidity to meet day to day liquid needs. Of course, banks are interested in maintaining as small a liquidity reserve as possible, because each unit of currency reserved for liquidity carries with it an opportunity cost owing to its uninvested nature, just as with a vanilla resource.

The total liabilities of a bank can again be described as a birth and death process on the set of non-negative integers, with the transition rates

$$\begin{array}{rcll} q(j,j+1) & = & \lambda, & j = 0,1,2,\ldots \\ q(j,j-1) & = & j\mu, & j = 1,2,3,\ldots \end{array}, \tag{6}$$

where λ is the rate of the deposit arrival process and μ is the rate of the withdrawal arrival process. Here state j refers to the bank holding j units of liability. Across the unbounded set of non-negative integers, the equilibrium distribution is (classically) given by the Poisson distribution (again, $\nu = \lambda/\mu$),

$$\pi_{Po}(j) = \frac{\nu^j}{j!}\left(\sum_{i=0}^{\infty}\frac{\nu^i}{i!}\right)^{-1} = \frac{\nu^j}{j!}e^{-\nu}, \qquad j = 0,1,2,\ldots. \tag{7}$$

This process becomes modified somewhat when $R-1$ resources are invested, and thereby removed from the system. In this case, the state space of the birth and death process is truncated *from below* and the transition rates become

$$\begin{array}{rcll} q(j,j+1) & = & \lambda, & j = R, R+1, \ldots \\ q(j,j-1) & = & j\mu, & j = R+1, R+2, \ldots \end{array} . \tag{8}$$

For this process, the equilibrium probability $\pi^l_R(R)$ (superscript l for lower truncated, and subscript R to again emphasize that different resource reservations lead to different equilibrium distributions) of occupying state R is then the boundary condition for normalization.

Through detailed balance, we derive the equilibrium distribution as

$$\pi^l_R(j) \quad = \quad \frac{\nu^j}{j!}\frac{R!}{\nu^R}\pi^l_R(R) \qquad\qquad j = R+1, R+2, \ldots, \tag{9}$$

and normalization provides $\pi^l_R(R)$,

$$\pi^l_R(R) \quad = \quad \frac{\nu^R}{R!}\left(e^\nu - \sum_{i=0}^{R-1}\frac{\nu^i}{i!}\right)^{-1} . \tag{10}$$

The final form of the equilibrium distribution then becomes

$$\pi^l_R(j) \quad = \quad \frac{\nu^j}{j!}\left(e^\nu - \sum_{i=0}^{R-1}\frac{\nu^i}{i!}\right)^{-1} \qquad\qquad j = R, R+1, \ldots . \tag{11}$$

If we compare this result to the equilibrium distribution of the birth and death process truncated from above in (2), we notice a disagreement. As an aside, however, notice that $\pi^l_0 = \pi_\infty = \pi_{Po}$.

The blocking probability of this system (arbitrarily denoted $F(\nu, R)$) is the fraction of withdrawals that must be handled via external means. By the reversibility of the underlying stationary linear open migration process, the withdrawal process at equilibrium is a Poisson process, and so the proportion of withdrawals being blocked is equal to the probability of being in the boundary state ($\pi^l_R(R)$),

$$F(\nu, R) = \pi^l_R(R) = \frac{\nu^R}{R!}\left(e^\nu - \sum_{i=0}^{R-1}\frac{\nu^i}{i!}\right)^{-1} . \tag{12}$$

This is then the fraction of withdrawals that the bank is not able to honour directly should they have $R-1$ units of currency occupied in investments, which is to be compared to $E(\nu, C)$ in (3), the fraction of resource subscription requests that a resource manager is forced to satisfy externally when only holding C units of the resource.

The differences between $F(\nu, R)$ and $E(\nu, C)$ will now be discussed. The two blocking probabilities pertain to opposing tail probabilities of a distribution that is very closely symmetric about its mean, provided that the traffic rate ν is large. And so, when we are interested in studying the 'lower blocking probability' $F(\nu, R)$, in practice it is much easier to study the 'upper blocking probability' $E(\nu, C)$. While we do not derive any limit theorems comparing the processes, we propose the following approximation based on symmetry (where $[\,\cdot\,]$ is the nearest integer function),

$$F(\nu, R) \approx E(\nu, 2[\nu] - R), \tag{13}$$

where the probability of a bank withdrawal denial is approximated by the blocking probability of a vanilla resource management problem with an appropriate capacity ($C = 2[\nu] - R$). Likewise, a vanilla resource management problem with capacity C approximately corresponds to a bank liquidity management problem with investment $R = 2[\nu] - C$.

By making this approximation, we are able to bring the problem of liquidity management into the standard language of resource management. The validity of the approximation will be motivated in Section 2.4.

2.2 On reserve requirements

Several national governments require banks to maintain a *reserve requirement* precisely for the liquidity management needs outlined above. From a historical perspective, reserve requirements were implemented to avoid *bank runs*, where banking customers questioning the solvency of a bank would rapidly withdraw funds. Reserve requirement make it possible for banks to manage irrational bank runs without being forced to liquidate long assets, a feedback amplification that could otherwise cause a major crisis for a bank.

In the United States the reserve requirement for banks holding more than \$43.9 million is 10% of deposits [2]. For banks in the United Kingdom there is no reserve requirement. In this essay, we will not be considering bank runs because they are principally a historical problem largely eliminated by deposit insurance schemes such as the *Federal Deposit Insurance Corporation* (FDIC) in the United States and *Financial Services Compensation Scheme* (FSCS) in the United Kingdom.

Our modeling framework is fully applicable when considering the management of liabilities under thinned deposit and withdrawal processes, where a certain percentage of incoming deposits is thinned off upon deposit and returned upon withdrawal. Therefore, while we have not explicitly considered reserve requirements, they are not in conflict with our approach.

2.3 Liquidity cost minimization

The operational challenge for banks is to determine the optimal amount of resources ($R - 1$) that should be invested for returns without the bank being forced to raise emergency capital through asset sales (a costly predicament, as mentioned previously). Such a 'fire sale' of assets is indeed problematic, and we will suppose the existence of an open money market, which is notably different from the interbank lending markets studied later in the essay.

On the open money market, banks can borrow money from money brokers outside the central banking system, and thereby avoid hazardous asset sales. While the interbank lending market closely tracks a target interest rate set by the central bank, open money market lending is much more obfuscated. The interest rates offered on the open money market fluctuate considerably. The British Bankers' Association conducts and releases daily industry polls tracking interest rates for loans in ten currencies at 15 loan durations (from overnight to 12 months). The interquantile means of these polls make up what are commonly called the LIBOR (*London InterBank Offered Rate*) rates [3].

Because the LIBOR lending market can at times be illiquid, it is important to consider ω_{EXT} as the general cost of acquiring liquidity 'externally', where a dried up LIBOR market may very well result in the fire sale of illiquid assets being the only accessible option.

Liquidity cost minimization is now a matter of balancing the profit from investing liabilities with the cost brought by being forced to perform external lending for unsatisfied withdrawals. For each unit of liquidity invested and removed from the bank, there is an investment profit (ω_{INV}), aloing with a probability $F(\nu, R)$ of incurring emergency external costs (ω_{EXT}). The optimization problem becomes

$$\begin{aligned} R_{opt}(\nu) &= \operatorname*{argmax}_{R \in \mathbb{Z}^+} \; (R-1)(\omega_{INV} - \omega_{EXT} F(\nu, R)), & (14) \\ &= \operatorname*{argmax}_{R \in \mathbb{Z}^+} \; \underbrace{(R-1)(1 - \frac{\omega_{EXT}}{\omega_{INV}} F(\nu, R))}_{\mathcal{C}_2(R)}. & (15) \end{aligned}$$

2.4 Approximation

To motivate the approximation in (13), we now compare the optima of the two cost functions $\mathcal{C}_1(C)$ and $\mathcal{C}_2(R)$ from (5) and (15). In Figure 2, $R_{opt}(\nu)$ is plotted alongside the approximation $2[\nu] - C_{opt}(\nu)$, with very good agreement. The interest rate scenario used is the last of the three presented in Figures 1 and 2. Furthermore, to show that this agreement is not specific to the chosen interest rate scenario, in Figure 3 we plot R_{opt} and $2[\nu] - C_{opt}$ for varying ω_{EXT} at a fixed traffic intensity $\nu = 100$, where we again find very good agreement.

In conclusion, we can approximate the lower blocking probability of the liquidity management problem by the upper blocking probability of the vanilla resource management problem. From this point in the essay and on, we will leave the details of liquidity management behind us and treat units of currency as any other vanilla resource, advancing to our main analysis of the effects of different network resource management strategies upon systemic risks.

2.5 Traffic burstiness

Throughout the essay, there is assumed to be a 'unit' of resource lending, which is an admittedly inaccurate assumption for liquidity management, where loans are certainly issued across a wide range of magnitudes, exhibiting notable burstiness. For this essay, it was thought best to confine the results to the Poissionian regime of traffic with uniform 'unit' sizes, and relegate possible generalizations for bursty traffic to future work.

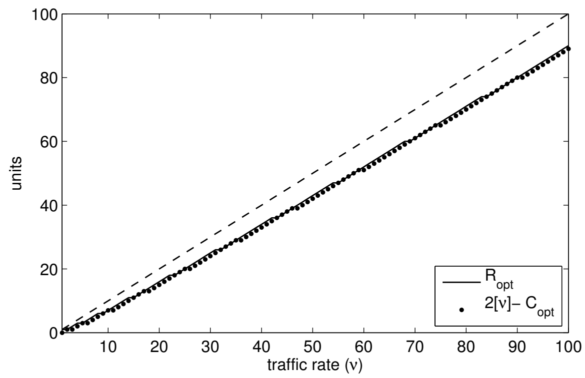

Figure 2: A comparison of the true optimal investment level of the liquidity problem (R_{opt}) with the approximation via the vanilla resource management problem ($2[\nu] - C_{opt}$), as a function of traffic intensity ν. The interest rate scenario is ($\omega_{EXT} = 0.09$, $\omega_{INV} = 0.03$).

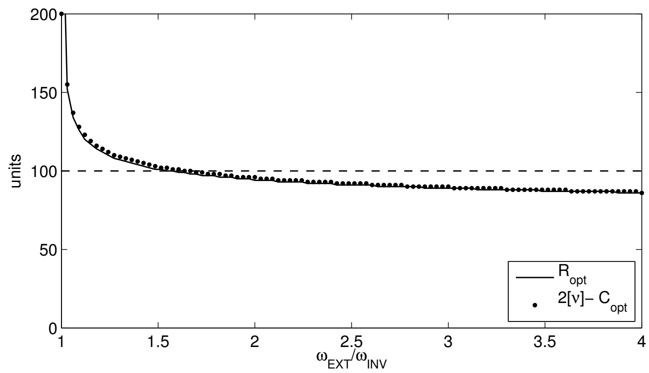

Figure 3: A comparison of the true optimal investment level of the liquidity problem (R_{opt}) with the approximation via the vanilla resource management problem ($2[\nu] - C_{opt}$), as a function of external investment cost ω_{EXT}, with fixed $\nu = 100$. Notice the saturation of the approximation at $\omega_{EXT}/\omega_{INV} = 1$, where $C_{opt} = 0$ implies that $2[\nu] - C_{opt} = 2[\nu] = 200$, while $R_{opt} \to \infty$. This inaccuracy is however well confined to the region very near $\omega_{EXT}/\omega_{INV} = 1$, and will not effect our modelling.

3 Lending networks

We are now prepared to study the consequences of interbank lending upon liquidity management. In practice, banks often borrow from each other to meet their liquidity needs (or rather, their reserve requirement needs, see Section 2.2). This occurs through the lending of *federal funds*, reserves kept with central banks specifically for this purpose, which is conducted at or near the *target federal funds rate*, a control parameter which has become the cornerstone of modern monetary policy.

There are many technical details regarding the terms of participation in the central bank lending network that we will not be elaborating in any detail. Notably, the existence of *discount window lending* plays an important role by offering overnight credit to participants, though excessive participation in the discount window lending process is discouraged through various mechanisms. In this essay, we assume that banks must turn the open money market (LIBOR) for needs not satisfied by the federal funds interbank market, rather than accessing the discount window. As a motivation for this decision, at the time of writing this essay the LIBOR rate is far above the discount window rate, indicating that banks are currently restricted from passing on discount window funds to the open money market, and the availability of discount window funds to cover liquidity needs is far from complete.

At this point it is important to emphasize that it is not the goal of this model to achieve completely detailed accuracy, but rather to investigate and illustrate the general structure of lending networks.

3.1 Network topology and symmetry

Throughout this essay, we assume a homogeneous banking network, where each bank is managing the same traffic rate ν. Further, to sidestep game-theoretic considerations of non-cooperation, we will assume each bank operates under the same management strategy, choosing the same liquidity capacity. We comment briefly on network heterogeneity and applications of game theory in Section 5.

Without interbank lending, each bank manages its liquidity independently, and as such, the trivial interbank network can be viewed as a directed star graph that we will call a *Trivial Lending Network (TLN)*. In a TLN, each bank is connected to a central *resource node (RN)* by a *resource link*, with capacity C_1, and banks independently manage withdrawal and deposit traffic arriving upon the directed edge from their bank to the RN, as in Figure 4(a). This is the 'network' studied in Section 1.

Building upon the Trivial Lending Network model, we now introduce two opposing directed edges between each pair of banks, which we call *interbank links*, enabling the possibility of routing withdrawal requests via other banks. We assign a homogenous capacity constraint (C_2) to these interbank links. We will call this extended graph of N bank nodes and one central resource node a *Interbank Lending Network* (ILN), see Figure 4(b). Momentarily disregarding the resource node, we see that the the bank nodes now form a complete directed graph.

The Interbank Lending Network model is rather different from the standard Loss Network model for telecommunications, as presented in [4]. That said, while the Interbank Lending Network model is not a complete graph, owing to

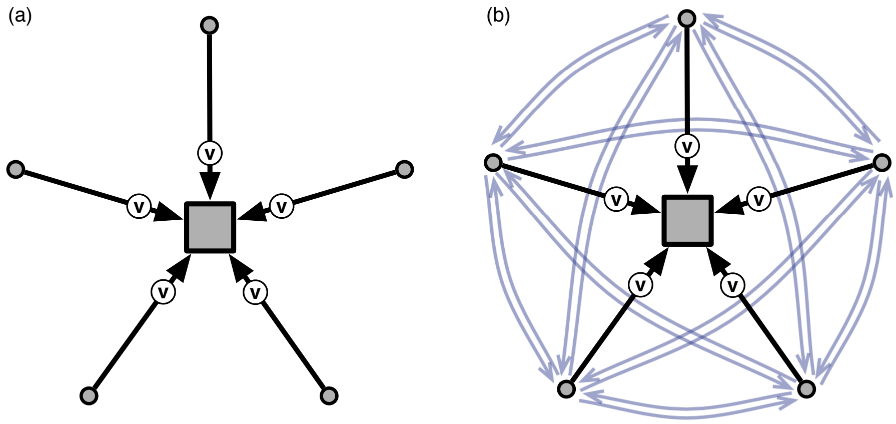

Figure 4: (a) The directed star graph representing a Trivial Lending Network with five independent banks and the resource node (RN) at the center. Traffic arrives independently at rate ν to each resource link connecting a bank to the RN. (b) The directed graph representing an Interbank Lending Network, with five banks and the RN at the center. It is important to note that no traffic arrives directly to the interbank links.

the resource node, it still exhibits notable symmetries, making the Erlang Fixed Point analysis commonly utilized in the study of Loss Networks fully applicable, as we will now show.

3.2 Withdrawal rerouting

When a withdrawal request arrives to a resource link between a bank and the resource node, the bank accepts the request if there is capacity. If there is no capacity available, the bank attempts to reroute the request via another bank, chosen at random (uniformly). The bank handling the rerouting request accepts the request if it has capacity.

Under this scheme, the traffic offered to each resource link is the regular arrival traffic for that bank (ν), as well as being offered a uniformly divided share ($1/(N-1)$) of the blocked traffic from $(N-1)$ other banks attempting to reroute. This traffic has been blocked once at the other bank's resource link with probability B_1 and then accepted on both the interbank link and the receiving resource link with probability $(1-B_1)(1-B_2)$.

Formalizing this rerouting mechanism, the probability that a withdrawal is blocked can be derived from the following Erlang Fixed Point calculation,

$$\left\{ \begin{array}{lcl} B_1 & = & E\left(\nu + \frac{\nu B_1(1-B_1)(1-B_2)}{(1-B_1)}, C_1\right) \\ B_2 & = & E\left(\frac{\nu B_1(1-B_1)(1-B_2)}{(1-B_2)}, C_2\right) \end{array} \right. , \tag{16}$$

which reduces to

$$\left\{ \begin{array}{lcl} B_1 & = & E\left(\nu + \nu B_1(1-B_2), C_1\right) \\ B_2 & = & E\left(\nu B_1(1-B_1), C_2\right) \end{array} \right. . \tag{17}$$

We advance immediately to the generalized scenario whereby M rerouting attempts are permitted. In this case matters do not simply so easily, and (16) becomes

$$\begin{cases} B_1 &= E\left(\nu + \frac{\nu B_1\left[1-[1-(1-B_1)(1-B_2)]^M\right]}{(1-B_1)}, C_1\right) \\ B_2 &= E\left(\frac{\nu B_1\left[1-[1-(1-B_1)(1-B_2)]^M\right]}{(1-B_2)}, C_2\right) \end{cases} . \tag{18}$$

Given a traffic rate ν, a set number of rerouting attempts M, and a set of capacities (C_1, C_2) to define our ILN, we can then calculate the fixed point (B_1^*, B_2^*) of the above equations. From this, we can calculate the probability that a withdrawal will be accepted for the ILN all-together as the sum of the probability of not being blocked on our first attempt $(1 - B_1^*)$ and the product of the probability of being blocked on the resource link on our first attempt and not being blocked for M rerouting attempts. This allows us to derive the *loss probability* L as

$$1 - L = (1 - B_1^*) + B_1^*\left(1 - [1 - (1 - B_1^*)(1 - B_2^*)]^M\right), \tag{19}$$

$$L = B_1^*\left[1 - (1 - B_1^*)(1 - B_2^*)\right]^M. \tag{20}$$

Note that this loss probability even applies when $M = 0$.

Here we briefly comment of the fact that the case $M = 0$ (the absence of routing) is equivalent to $C_2 = 0$. In later sections, we will primarily be considering only the case of aggressive loan rerouting, $M = 10$, and characterize the lending market (or lack there of) solely through the interbank capacity C_2. The effect of varying M is returned to in Section 4.1.

3.3 Interbank link capacity (C_2), extreme cases

Let us consider the extreme cases of interbank lending, which take on notable reduced forms that will return later in the essay. If the interbank capacity is infinite, $C_2 = \infty$, this leads to $B_2 = 1$, and (18) and (20) reduce to

$$B_1 = E\left(\nu + \frac{\nu B_1\left[1 - B_1^M\right]}{(1 - B_1)}, C_1\right), \tag{21}$$

$$L' = [B_1^*]^{M+1}. \tag{22}$$

At the other extreme, let us consider what happens when the interbank capacity is to $C_2 = 0$. In this case $B_2 = 1$, and (18) reduces to the Trivial Lending Network scenario studied in Section 1, where we obtain the explicit Erlang formulas

$$B_1^* = E(\nu, C_1), \tag{23}$$

$$L'' = B_1^*. \tag{24}$$

3.4 The limit $\nu \to \infty$, $C_1 \to \infty$

In this subsection, we derive limit results for the blocking probability and loss probability under the limiting regime $\nu \to \infty$, $C_1 \to \infty$, where ν/C_1 is constant, analogous to the standard result presented for independent links and applied to Loss Networks in [4]. We conjecture that the limit form of the loss probability L is independent of the interbank capacity C_2 and the number of rerouting attempts M. Furthermore, we propose that this implies that Interbank Lending Networks are never hysteretic, in contrast to the occasional bistability of Loss Networks under certain high traffic conditions.

We first review the standard limit result for independent lending networks ($C_2 = 0$, $M = 0$), which is the same as for an independent link in a Loss Network.

Lemma 3.1. *For an independent link with traffic rate λ and capacity C, the high-traffic-high-capacity limit of the loss probability is $L = (1 - \frac{C}{\lambda})^+$.*

Proof. We are interested in studying what happens to $E(N\lambda, NC)$ as $N \to \infty$. We find the limiting blocking probability B to be

$$\begin{aligned}
B &= \lim_{N\to\infty} E(N\lambda, NC) && (25)\\
&= \lim_{N\to\infty} \frac{\frac{(N\lambda)^{NC}}{(NC)!}}{\sum_{i=0}^{NC} \frac{(N\lambda)^j}{j!}} && (26)\\
&= \lim_{N\to\infty} \frac{1}{1 + \frac{C}{\lambda} + \cdots + \left(\frac{C}{\lambda}\right)^{NC}} && (27)\\
&= \left(1 - \frac{C}{\lambda}\right) \lim_{N\to\infty} \frac{1}{1 - \left(\frac{C}{\lambda}\right)^{NC+1}} && (28)\\
&= \begin{cases} 1 - \frac{C}{\lambda} & \text{if} \quad C < \lambda \\ 0 & \text{if} \quad C > \lambda \end{cases} && (29)\\
&= \left(1 - \frac{C}{\lambda}\right)^+ && (30)
\end{aligned}$$

For independent links, the loss probability $L = B$, which finishes the result. ☐

Next, we show that the same loss probability limit is achieved for $M = 1$, which is importantly independent of C_2.

Lemma 3.2. *For an Interbank Lending Network with rerouting count $M = 1$, traffic rate ν, resource link capacity $C_1 \geq 0$ and interbank link capacity $C_2 \geq 0$, the high-traffic-high-capacity limit of the loss probability of the network is $L = (1 - \frac{C_1}{\nu})^+$.*

Proof. Applying the limit from Lemma 3.1 to (17), we obtain

$$\left\{ \begin{aligned} B_1 &= \left(1 - \frac{C_1}{\nu + \nu B_1(1-B_2)}\right)^+ \\ B_2 &= \left(1 - \frac{C_2}{\nu B_1(1-B_1)}\right)^+ \end{aligned} \right. . \qquad (31)$$

We proceed with two cases. In the first case, $C_2/\nu > B_1(1 - B_1)$, which implies that $B_2 = 0$, and the blocking and loss probabilities become

$$\begin{aligned}
B_1^2 &= \begin{cases} 1 - \frac{C_1}{\nu} & \text{if} \quad C_1 < \nu \\ 0 & \text{if} \quad C_1 > \nu \end{cases} , && (32)\\
L &= B_1(1 - (1 - B_1)) = B_1^2 && (33)\\
&= \left(1 - \frac{C_1}{\nu}\right)^+ . && (34)
\end{aligned}$$

In the second case, $C_2/\nu < B_1(1 - B_1)$, we find a different blocking probability, but the same loss probability,

$$
\begin{aligned}
B_1 &= \begin{cases} 1 - \frac{C_1}{\nu} + \frac{C_2}{\nu} & \text{if} \quad C_1 < \nu \\ 0 & \text{if} \quad C_1 > \nu \end{cases}, && (35) \\
L &= B_1(1 - (1 - B_1)(1 - B_2)) && (36) \\
&= \begin{cases} (1 - \frac{C_1}{\nu} + \frac{C_2}{\nu}) - \frac{C_2}{\nu} & \text{if} \quad C_1 < \nu \\ 0 & \text{if} \quad C_1 > \nu \end{cases} && (37) \\
&= \left(1 - \frac{C_1}{\nu}\right)^+. && (38)
\end{aligned}
$$

□

Now, we conjecture that the same holds for any general M.

Conjecture 3.3. *For an Interbank Lending Network with traffic rate ν and resource link capacity $C_1 \geq 0$, the high-traffic-high-capacity limit of the loss probability of the network is $L = (1 - \frac{C_1}{\nu})^+$, for all $M \geq 0$, $C_2 \geq 0$.*

In addition to the lemma proven above, the conjecture is further made plausible by another lemma, showing that the conjecture is provably true for all choices of M in the special case when $C_2 = \infty$.

Lemma 3.4. *For an Interbank Lending Network with traffic rate ν, resource link capacity $C_1 \geq 0$ and infinite interbank link capacity $C_2 = \infty$, the high-traffic-high-capacity limit of the loss probability of the network is $L = (1 - \frac{C_1}{\nu})^+$, for all $M \geq 0$.*

Proof. If we apply Lemma 3.1 to the M-try rerouting strategy with infinite C_2 presented in (21), we obtain

$$
\begin{aligned}
B_1 &= \left(1 - \frac{C_1(1 - B_1)}{\nu(1 - B_1) + \nu B_1(1 - B_1^M)}\right)^+ && (39) \\
&= \left(1 - \frac{C_1(1 - B_1)}{\nu(1 - B_1^{M+1})}\right)^+. && (40)
\end{aligned}
$$

After some algebra, we can derive

$$
B_1^{M+1} = \left(1 - \frac{C_1}{\nu}\right)^+, \tag{41}
$$

and since the loss probability $L = B_1^{M+1}$, we again obtain the same familiar result as before, $L = (1 - \frac{C_1}{\nu})^+$. □

This makes us fairly confident that our conjecture is true. What is the significance of this conjecture, if true? In the case of Loss Networks, traffic rerouting is known to causes hysteresis effects for high traffic systems (with accordingly high capacity) [5]. In Lemma 3.4, we have effectively shown that the high-traffic-high-capacity limit of the loss probability for a lending network with infinite interbank lending capacity C_2 is not hysteretic, no matter how aggressively traffic is rerouted.

This is an important difference between Loss Networks and ILNs. Since ILNs do not exhibit hysteresis, technical strategies such as *trunk reservation* [4], which can be shown to resolve hystereses for Loss Networks, are not necessary. Intuitively this is sensible, since for ILNs, the interbank links utilized for rerouting are not in competition for direct routing. Having motivated this non-existance of hysteresis through the above limit results, we now return to the more concrete operational task of cost minimization.

3.5 Cost minimization and rerouting efficiency

We now investigate what happens when the capacity constraints C_1 and C_2 are chosen optimally, though a cost-minimization procedure analogous to the approaches used in Section 1 and 2. Will once again limit ourselves to the case where all banks receive the same traffic rate ν and all banks choose the same operating capacities.

This symmetry leads to a very important degeneracy in our model. Because all banks have the same capacity strategy, traffic, and interest rates, the expected frequency of lending is equal to the expected frequency of borrowing. Because of this, there is no cost of lending. As such, the lending network will exhibit high systemic connectivity, a property we will explore further in the next section, where this symmetry will be broken by a probability of loan default.

Without a cost of lending or borrowing, the operational cost function then consists of the sum of the cost of capital ($\omega_{INV}C$) and the cost of external resource acquisition for lost withdrawal requests ($\omega_{EXT}\nu L$), though the loss probability L for the network problem, as derived in (20), is notably different from the earlier independent problems, where the loss probability was simply the blocking probability. The optimization problem becomes

$$\begin{aligned} \mathbf{C}_{opt} &= \underset{C_1,C_2\in\mathbb{Z}^+}{\operatorname{argmin}}\ \omega_{INV}C_2 + \omega_{EXT}\nu L && (42)\\ &= \underset{C_1,C_2\in\mathbb{Z}^+}{\operatorname{argmin}}\ C_1 + \frac{\omega_{EXT}}{\omega_{INV}}\nu\left[B_1^*\left(1-(1-B_1^*)(1-B_2^*)\right)^M\right], && (43)\end{aligned}$$

where (B_1^*, B_2^*) are the fixed point to the multirouting Erlang equations,

$$\begin{cases} B_1 = E\left(\nu + \frac{\nu B_1\left[1-[1-(1-B_1)(1-B_2)]^M\right]}{(1-B_1)}, C_1\right) \\ B_2 = E\left(\frac{\nu B_1\left[1-[1-(1-B_1)(1-B_2)]^M\right]}{(1-B_2)}, C_2\right) \end{cases}, \qquad (44)$$

and the optimum $\mathbf{C}_{opt}$ is now vector valued. There is no direct cost associated with the interbank link capacity C_2, since these links do not correspond to actual resources, which differs from the resource link capacity C_1.

In Figure 5(a), we compare this cost function to the cost function for the independent management problem described in Section 1. In Figure 5(b), the optimal resource link capacities $C_{1,opt}$ is compared to that of the independent banking problem, as a function of the traffic intensity ν, with $M = 10$ fixed. Not shown is the fact that in all these studies, $C_{2,opt} = \infty$, since there are no negative consequences of lending or borrowing.

Interbank lending clearly decreases the operational cost of liquidity management, as the considerable difference between the minimal operational costs shown in Figure 5(a) confirm.

(a)

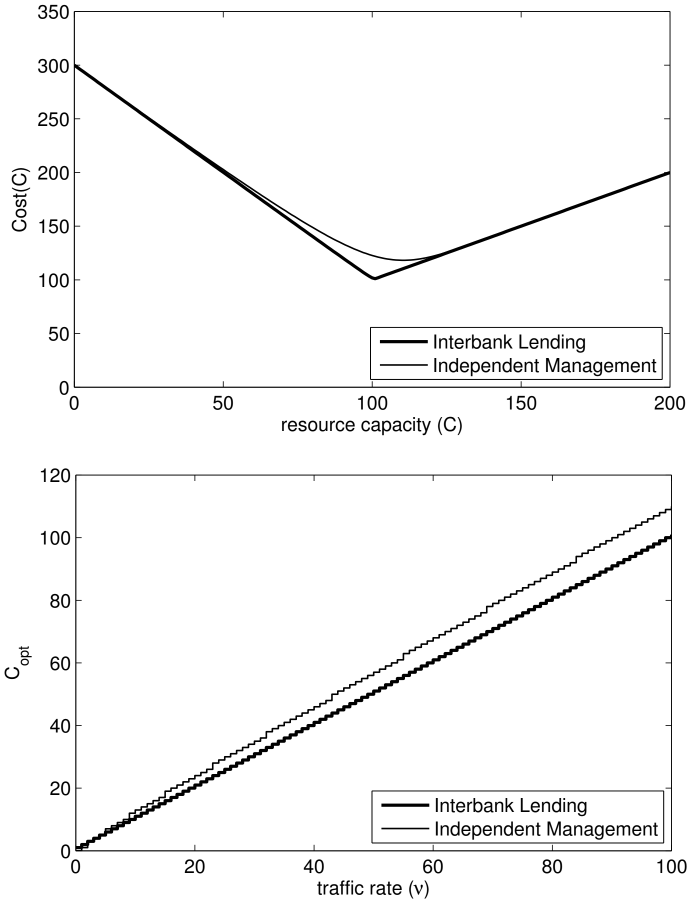

(b)

Figure 5: (a) The cost function for the independent resource management problem from (5), compared to the cost function for the interbank lending problem ($M = 10$) from (43), at traffic rate $\nu = 100$. (b) The optimal resource link capacity for the independent management problem C_{opt}, compared to the optimal resource link capacity $C_{1,opt}$ for the network problem ($M = 10$), where the optimal interbank link capacity $C_{2,opt} = \infty$, for all ν. The interest rate scenario here is ($\omega_{INV} = 0.03, \omega_{EXT} = 0.09$).

4 Systemic connectivity and default risks

In the previous section we claimed that since the expected profit of lending and the expected cost of borrowing are equal and opposite in sign for participants in an Interbank Lending Network with symmetric traffic, there is no cost associated with lending participation. We will now show what lending does to the systemic connectivity of the network, and then introduce an asymmetric cost of loan defaults, with curious results.

4.1 Systemic connectivity

Here we study the fraction of withdrawal requests that are rerouted when interbank capacity is enabled ($C_2 = \infty$), as a function of the resource capacity C_1, paying particular attention to this rerouting fraction for the optimal capacity $C_{1,opt}$ from the optimization problem in the previous section.

The probability that a withdrawal is rerouted is the probability that it is blocked times the probability that it is not blocked in one of M rerouting attempts. This probability, divided by the probability that a request is not lost, gives us the fraction of carried requests that are rerouted, which we will utilize as a measure of systemic connectivity. This fraction R is given by

$$R = \frac{B_1^*\left(1-[1-(1-B_1^*)(1-B_2^*)]^M\right)}{1-L} \tag{45}$$

$$= \frac{B_1^*\left(1-[1-(1-B_1^*)(1-B_2^*)]^M\right)}{(1-B_1^*)+B_1^*\left(1-[1-(1-B_1^*)(1-B_2^*)]^M\right)} \tag{46}$$

$$= \frac{B_1^*-B_1^*\,[1-(1-B_1^*)(1-B_2^*)]^M}{1-B_1^*\,[1-(1-B_1^*)(1-B_2^*)]^M}. \tag{47}$$

When $C_2 = \infty$, as is the case for the optimal management strategy in the previous section, then $B_2^* = 0$ and this reduces to

$$R' = \frac{B_1^* - B_1^{*M+1}}{1-B_1^{*M+1}}. \tag{48}$$

In Figure 6 we plot the rerouting fraction R as a function of resource capacity C_1, for traffic intensity $\nu = 100$, at several different rerouting attempt levels M. We highlight the location of the optimal capacity $C_{1,opt}$ in each case, and alongside R we also plot the loss probabilities L and blocking probabilities B_1^*. All plots utilize the interest rate scenario ($\omega_{INV} = 0.03, \omega_{EXT} = 0.09$).

In Figure 7, we offer a more pointed view of the effect of varying the number of rerouting attempts M. We plot the same three quantities (B_1^*, L and R) at the optimal resource link capacity $C_{1,opt}$, as a function of rerouting attempts M. We also plot the optimum $C_{1,opt}$ itself, along with the cost at the optimum. Here we can see that beyond $M = 10$, the cost benefit of rerouting is no longer measurable, while the rerouting fraction R, our measure of connectivity, continues to rise.

As $M \to \infty$, we see the rerouting fraction R approach 1. This is the systemic connectivity which we are referencing. While aggressive rerouting does increase the economic efficiency of the network, the fraction of withdrawals that are being handled in a rerouted manner increases sharply, to a point where any given bank is handling almost entirely withdrawals from other banks. Next we will see what happens when we break the symmetry of the lending benefit.

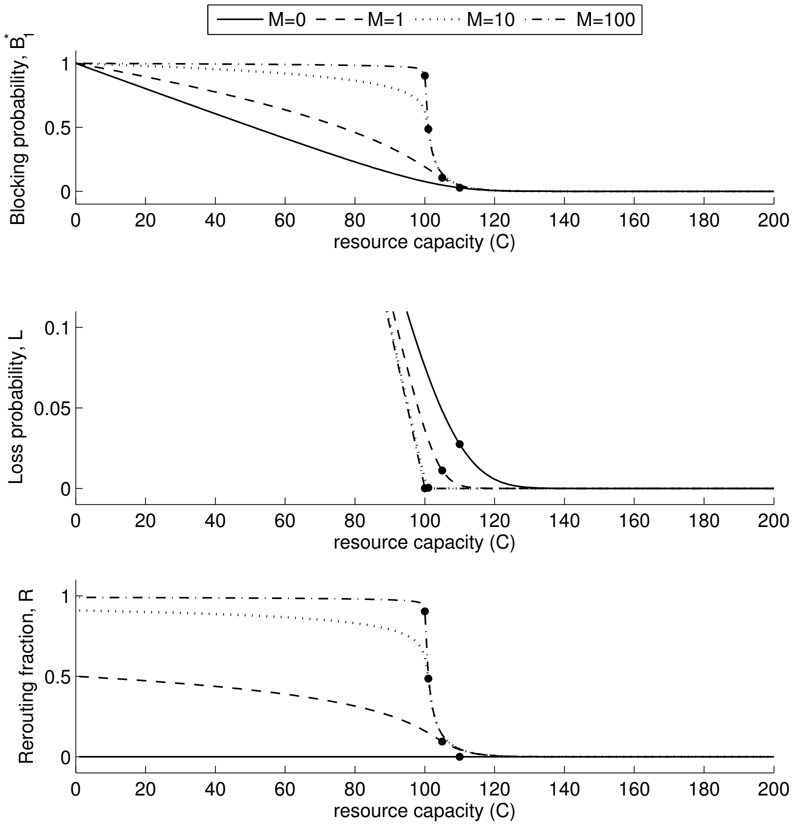

Figure 6: The blocking probability B_1^*, loss probability L and rerouting fraction R as a function of resource link capacity, for $\nu = 100$ and four different rerouting intensities, $M = 0$, 1, 10, and 100. The optimal resource link capacities $C_{1,opt}$ are indicated by dots. As rerouting intensity is increased, the fraction of withdrawals that each bank is rerouting through other banks grows considerably, creating systemic connectivity. The loss probability is plotted on a different vertical scale in order to illustrate the relevant region of interest. All plots utilize the interest rate scenario $(\omega_{INV} = 0.03, \omega_{EXT} = 0.09)$.

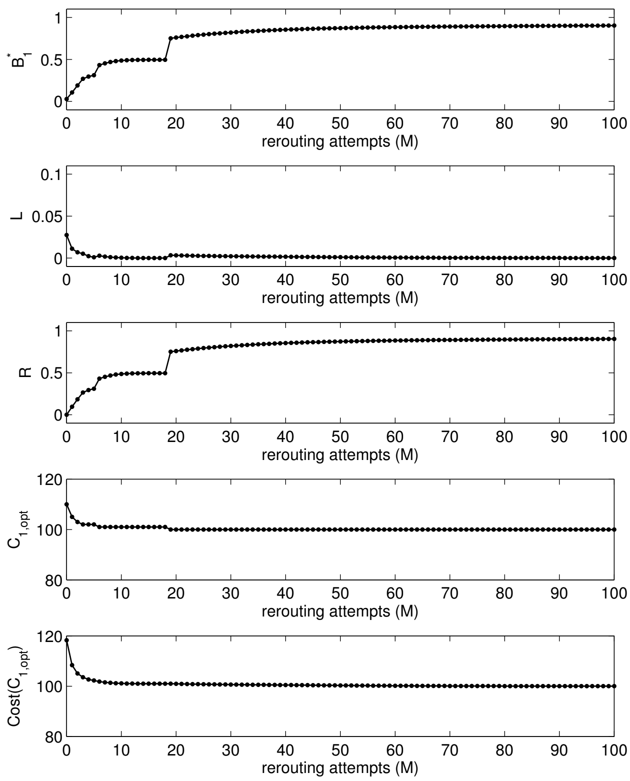

Figure 7: The blocking probability B_1^*, loss probability L and rerouting fraction R at the optimal resource capacity $C_{1,opt}$, in addition to $C_{1,opt}$ itself and the minimal cost, all as a function of varying the rerouting attempts count M. The discontinuous jumps in B_1^* and R correspond to downward shifts in the optimal resource link capacity. Notice that beyond $M = 10$, the cost benefit of further rerouting is no longer significant, while the connectivity (R) continues to rise considerably. Again the traffic intensity is $\nu = 100$ and the interest rate scenario is $(\omega_{INV} = 0.03, \omega_{EXT} = 0.09)$.

4.2 Default risk

We now aim to modify the cost minimization problem from section 3.5 to explicitly incorporate a cost for loan defaults. In addition to the terms of the previous cost function, we add a loan default term given by the probability of default ρ times the fraction of all offered traffic that is routed via other banks,

$$\mathbf{C}_{opt} = \mathrm{argmin}_{C_1, C_2 \in \mathbb{Z}^+} \quad \begin{array}{l} \omega_{INV} C_1 + \omega_{EXT} \nu L \\ + \rho \nu B_1^* \left(1 - [1 - (1 - B_1^*)(1 - B_2^*)]^M\right) \end{array} \tag{49}$$

$$\mathbf{C}_{opt} = \mathrm{argmin}_{C_1, C_2 \in \mathbb{Z}^+} \quad \begin{array}{l} \omega_{INV} C_1 \\ + \omega_{EXT} \nu B_1^* \left[1 - (1 - B_1^*)(1 - B_2^*)\right]^M \\ + \rho \nu B_1^* \left(1 - [1 - (1 - B_1^*)(1 - B_2^*)]^M\right). \end{array} \tag{50}$$

As the default risk ρ is increased, we observe a rapid transition in optimal lending strategy similar to a phase transition from physical systems, where the benefit of lending participation transitions discontinuously to a detriment beyond a critical probability of default ρ^*. To examine this critical transition in detail, Figure 8 plots the marginal cost function over interbank capacities C_2 and default risks ρ, where the resource capacity C_1 has already been optimized out. The other parameter values are $\nu = 100$, $M = 10$, $\omega_{INV} = 0.03$, $\omega_{EXT} = 0.09$.

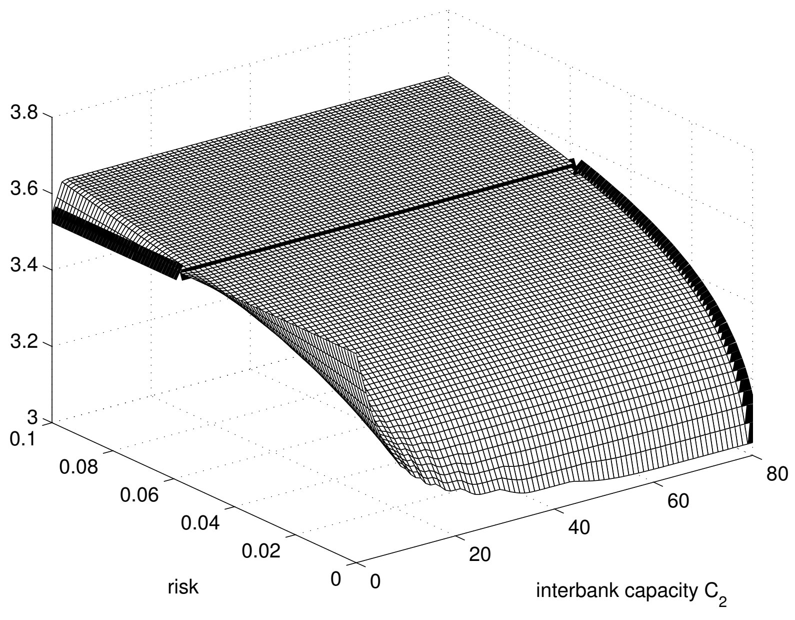

Figure 8: The marginal cost function for different interbank capacities C_2 and default risks ρ, where the choice of resource capacity C_1 has already been optimized out. The optimal $C_{2,opt}(\rho)$ is bolded, illustrating how $C_{2,opt}$ transitions abruptly from ∞ to 0. The other parameters are $\nu = 100$, $M = 10$, and the interest rate scenario is $(\omega_{INV} = 0.03, \omega_{EXT} = 0.09)$.

To more broadly characterize the critical transition, Figure 9 presents the effect of the default risk ρ upon the blocking probability B_1^*, loss probability L, rerouting fraction R, optimal capacities $\mathbf{C}_{opt} = (C_{1_opt,}, C_{2,opt})$, and the minimized cost. The calculations are again run for parameter values $\nu = 100$, $\omega_{INV} = 0.03$, $\omega_{EXT} = 0.09$, and $M = 10$. We see that under the interest rate/traffic scenario being studied, the phase transition at approximately $\rho^* = .059$ causes a discontinuous rise in the loss probability.

This transition implies that there is a critical probability of default risk at which rational participants in an optimized Interbank Lending Network will simultaneously stop lending, leading to a 'collapse' of the lending network. This would in turn rapidly push large amounts of loan traffic to the external money

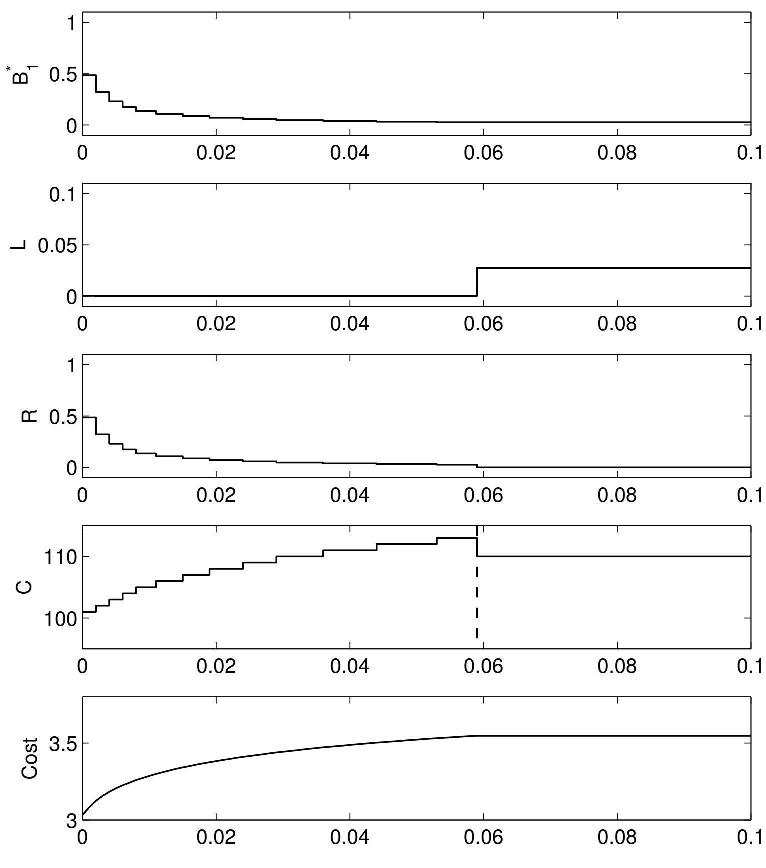

Figure 9: The blocking probability B_1^*, loss probability L, rerouting fraction R, optimal resource capacitiy $C_{1,opt}$, and minimized cost, as the probability of loan default, ρ, is varied. The critical transition point for $C_{2,opt}$ is shown (dashed) alongside $C_{1,opt}$. The traffic intensity is $\nu = 100$, the rerouting intensity is $M = 10$, and the interest rate scenario is $(\omega_{INV} = 0.03, \omega_{EXT} = 0.09)$.

market, and in a more accurate model, this would realistically increase the cost of external lending considerably (for the purposes of our discussion, LIBOR), and a widespread liquidity crisis would become plausible.

It is important to note that there are many unrealistic assumptions behind this model, in particular the symmetry of bank traffic. The purpose of the essay is not to offer a definitive or even notably accurate model of interbank lending, but rather to illustrate the problematic fragility of lending systems in general.

4.3 Critical default risk

Here we return to the significance of the interest rate parameters ω_{INV}, the costs of investment capital, and ω_{EXT}, the cost of external funds. Throughout sections 3 and 4, we have been operating under a single interest rate scenario, ($\omega_{INV} = 0.03$, $\omega_{EXT} = 0.09$). This scenario was chosen because it is a representative non-trivial choice that illustrates well the intricacies of the problem. We will now briefly show how the critical risk depends upon the interest rate parameters.

In Figure 10 we present the critical risk ρ^*, and its surprisingly simple dependence upon the two interest rates, calculated at $\nu = 100$ and $M = 10$. We have tried to derive a simple expression for ρ^* by studying the cost function in (50) and the difference between the optimal C_1 conditioned upon the competing scenarios $C_2 = 0$ and $C_2 = \infty$, but because this involves the optima of optimization problems evaluated over terms derived from fixed point calculations, no simple form has presented itself.

Also shown in Figure 10 is the change in loss probability when the phase transition occurs, also generally considered as a function of the interest rate parameters. We again emphasize that a sudden large increase in the external lending traffic would undoubtedly have an effect upon the interest rate for external funds, which clearly raises questions about the optimality of the 'optimal' behavior. This feedback mechanism complicates the procedure considerably, and we let it remain a subject for future study.

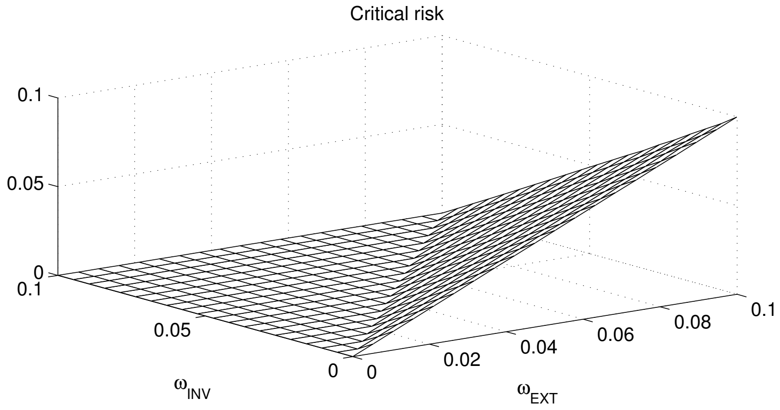

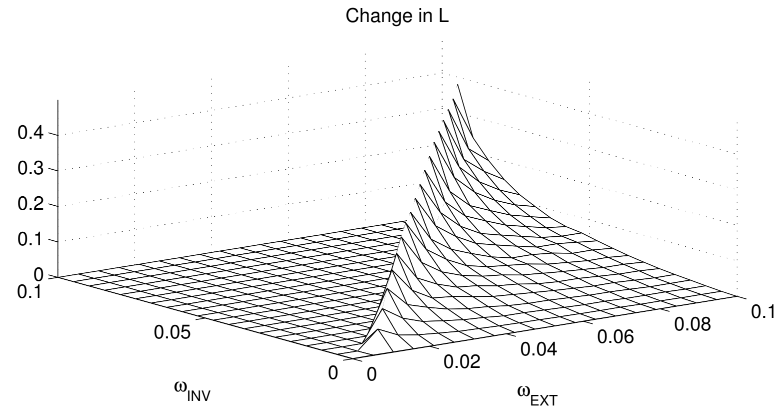

Figure 10: (Above) The critical value of the probability of default, ρ^*, plotted as a function of the alternative investment opportunity interest rate, ω_{INV}, and the interest rate for external funds, ω_{EXT}. (Below) The corresponding change in the loss probability L at the critical point. We see that when $\omega_{INV} > \omega_{EXT}$, we have no critical risk, since all traffic is being directed at the external money market. All calculations at $\nu = 100$ and $M = 10$.

5 Conclusions and future directions

In this essay we have presented a model of interbank lending developed using the mathematics of network resource management. We have used the model to study the effects of lending upon the efficiency of liquidity management, given a complete network of symmetric bank traffic. Furthermore, we have examined the emergence of systemic connectivity in the model, and noted how the risk of loan defaults creates a phase transition in optimal resource management strategy. We observe that a small change in default risk can therefore lead to a systemic liquidity crisis.

The assumptions of traffic symmetry is a key assumption that distances the results measurably from real banking systems [6, 7]. Techniques for analyzing non-symmetric traffic within the loss network framework exist [4], and absolutely merit application in future studies. Such an analysis has the potential to characterize the often-discussed question of what effect comparatively large banks have upon a lending market, potentially providing important insight into the matter of how and when banks become 'too big to fail'. Recent popular discussions have suggested that larger banks should perhaps be required to keep a larger percentage of their funds as reserve requirement. These suggestions seem loosely based upon an intuition that when a single bank is much larger than the other agents of the lending market, its liquidity also fluctuates on a different order of magnitude than the other agents, meaning that it can have a very difficult time obtaining loans at the order of magnitude it may periodically require, should it encounter a sizable drift in available liquidity. Therefore large banks may need to manage the volatility of their liquidity more independently, keeping more liquidity on hand than other smaller agents. The thoughts noted here are only speculation, but an investigation into 'too big to fail' scenarios is a natural extension of the work presented in this essay.

The restriction to cooperative resource management strategies, where all network agents make the same resource capacity decisions, is also notably different from the non-cooperative reality of lending markets. In reality, each bank decides on its own how much resource capacity to keep on hand, and to what extend it should accept or apply for interbank loans. This transforms the problem into a high-dimensional non-cooperative multiplayer game, which is more than we've been able to consider in this introductory study. A game theoretic consideration of the systems studied in this essay would make for very interesting future work.

Lastly, the model could also be refined to more closely describe the day-to-day operation of the lending market. In reality, central bank lending markets generously allow for intraday 'daylight overdraft', and banks are only forced to bring themselves into compliance with reserve requirements at the close of the market. This operational detail can lead to very complex intraday behavior [8], parts of which are game-theoretic in nature, and should not be disregarded.

As a closing acknowledgment, the 'efficiency-induced fragility' perspective of interbank lending networks presented in this essay is inspired by (though notably different in nature from) classic robustness conservation laws from control theory [9] and more recent results from systems biology [10, 11].

References

[1] E. Brockmeyer *et al.*, *The life and works of AK Erlang.* The Danish Academy of Technical Sciences, 1960.

[2] The Federal Reserve Board, http://www.federalreserve.gov/pubs/supplement/2008/02/, *Statistical Supplement to the Federal Reserve Bulletin*, February 2008.

[3] D. MacKenzie, "What's in a number? the importance of libor," *London Review of Books*, vol. 30, September 2008.

[4] F. Kelly, "Loss networks," *Annals of Applied Probability*, vol. 1, no. 3, pp. 319–378, 1991.

[5] R. Gibbens, P. J. Hunt, and F. P. Kelly, "Bistability in communication networks," in *Disorder in physical systems* (G. Grimmett and D. Welsh, eds.), pp. 113–127, Oxford University Press, 1990.

[6] K. Soramäki, M. Bech, J. Arnold, R. Glass, and W. Beyeler, "The topology of interbank payment flows," *Physica A: Statistical Mechanics and its Applications*, vol. 379, no. 1, pp. 317–333, 2007.

[7] M. Bech and E. Atalay, "The topology of the federal funds market," *Federal Reserve Bank of New York Staff Reports*, no. 354, 2008.

[8] M. Bech and R. Garratt, "Illiquidity in the interbank payment system following wide-scale disruptions," *Federal Reserve Bank of New York Staff Reports*, no. 239, 2006.

[9] H. W. Bode, *Network analysis and feedback amplifier design.* Van Nostrand Reinhold, Princeton, NJ, 1956.

[10] M. Csete and J. Doyle, "Reverse engineering of biological complexity," *Science*, vol. 295, no. 5560, pp. 1664—1669, 2002.

[11] H. Kitano, "Towards a theory of biological robustness," *Molecular Systems Biology*, vol. 3, no. 1, 2007.