# Financial Contagion in the Laboratory: Does Network Structure Matter?*


| John Duffy | Aikaterini Karadimitropoulou | Melanie Parravano |
|---|---|---|
| UC Irvine | University of East Anglia | Newcastle University |
| duffy@uci.edu | A.Karadimitropoulou@uea.ac.uk | Melanie.Parravano@ncl.ac.uk |


June 2016


**Abstract:** We design and report on laboratory experiments exploring the role of interbank network structure for the likelihood of a financial contagion. The laboratory provides us with the control necessary to precisely explore the role of different network configurations for the fragility of the financial system. Specifically, we study the likelihood of financial contagion in complete and incomplete networks of banks who are linked in terms of interbank deposits as in the model of Allen and Gale (2000). Subjects play the role of depositors who must decide whether or not to withdraw their funds from their bank. We find that financial contagions are possible under both network structures. While such contagions always occur under an incomplete interbank network structure, they are significantly less likely to occur under a complete interbank network structure where interbank linkages can effectively provide insurance against shocks to the system, and localize damage from the financial shock.



**Keywords**: Contagion, Networks, Experiments, Bank Runs, Interbank Deposits, Financial Fragility.

**JEL Codes**: C92, E44, G21.



*** Acknowledgments:** We wish to thank participants of presentations at the UEA Economics School seminar, the CBESS seminar (UEA) for useful advice and feedback on earlier drafts. We also wish to thank the Centre for Behavioural and Experimental Social Sciences (CBESS), UEA, and UC Irvine for funding and Sara Godoy for providing support to organize the experimental sessions. The usual disclaimer applies.

## 1. Introduction

The financial crisis of 2007-08 has reinforced the view that interbank network linkages are crucial to understanding the financial fragility of a country's banking system. Unlike earlier financial crises, the crisis of 2007-08 did not simply involve depositors running to withdraw money from their own banks. Rather, it also involved banks (and some large nonbanks) with interbank deposits running on other banks holding those deposits. For example, in the U.S., the collapse of Lehman Brothers was associated with a $423 billion dollar contraction in the U.S. dollar *interbank* lending market (Gorton 2010), and this in turn pushed other banks to the brink requiring government bailouts (e.g. Morgan Stanley) or led them to be sold off (e.g., Merrill Lynch).

The traditional view of financial crises as involving a run by bank depositors on their own bank has been modeled as a self-fulfilling equilibrium coordination game by Diamond and Dybvig (1983) where depositor's beliefs play a pivotal role. The more modern view of *financial contagion* as an equilibrium phenomena arising from the interbank *network structure* was first proposed by Allen and Gale (2000). While the Diamond and Dybvig model involves the behavior of depositors in a single bank, Allen and Gale's model considers the behavior of depositors across many, interconnected banks. In this paper we explore the key implication of Allen and Gale's interbank model of financial crises, namely that the network structure *matters* for the fragility of the banking system. We address the importance of network structure for financial fragility using the methodology of experimental economics, which provides us with precise control over the network structure of interbank connections as well as over the information that is available to depositors in that network. This control enables us to gather data that can be used to directly test the role played by network structure in the spread of a financial crisis. While there are many experimental studies of the Diamond and Dybvig model of bank runs our paper provides the first experimental test of whether the *interbank network structure* matters for the likelihood of financial contagion.

In our experiment, subjects are depositors in their local regional bank. As in the Allen and Gale model, this regional bank is one of four interconnected banks in the economy. Each bank holds deposits with other banks as a means of insuring against the uncertain liquidity demands of their depositors. Following Allen and Gale (2000) we focus on two different interbank market structures, namely, an "incomplete" market structure where the banks are partially connected (i.e., each bank holds deposits in one adjacent bank) and a "complete" market structure, where the four banks are fully connected (i.e., each bank invests a fraction of their deposits in each one of the other three banks). According to the model, as detailed in the next section, the introduction of interbank linkages (i.e., exchange of deposits) implies that both market structures can implement the first-best (i.e., no bank run equilibrium). However, the network structure can lead to important differences in response to an exogenous liquidity shock. Our experiment was designed to test the implications of such fragility and for this reason, in every round we introduce a liquidity shock to one of the four banks in the economy. Depositors' payoffs are carefully calibrated to capture the model assumptions. The result is a risk-sharing coordination game, where there is as a unique Nash equilibrium when the interbank network is incomplete involving full contagion. By contrast, when the interbank network structure is complete, both the inefficient and efficient equilibrium coexist allowing the possibility of either full contagion or no-contagion, respectively. We implement two treatments between-subjects, one for each 'network structure' and participants played 30 rounds of a game in which they were repeatedly confronted with the choice of withdrawing or keeping their deposits with their local regional bank. Our main research objective is to understand whether a more integrated and transparent banking system leads to smaller self-fulfilling spillover effects, as predicted by the model.

To preview our findings, our main experimental result is that a complete interbank network structure is indeed likely to reduce the risk of contagion. In particular, we find that under the *incomplete* interbank network structure all economies converged to an outcome approximating the full-contagion equilibrium whereas under the *complete* network structure only about half did. Econometric analysis also shows that under the incomplete network structure the probability of a participant withdrawing her deposit was significantly larger than in the complete network treatment, this even after controlling for the past behaviour of co-players and own past behaviour as well. Moreover, in the incomplete network structure, we also observe the expected *pattern of contagion*, where there is a spillover from the shocked bank to the bank directly connected with that bank and then to the next bank until finally the full banking network is affected. Therefore, our results provide support for the model's prediction that in an incomplete interbank network structure, an initial financial shock spreads to all banks and the crisis becomes global. We also provide partial support to the models' prediction that in a fully integrated banking system (a complete network structure) a financial crisis does not become global.

The rest of the paper is organized as follows. The next section situates our paper in the relevant literature. Section 3 presents the model and the main hypotheses concerning the consequences of network structure for financial contagion. Section 4 describes our experimental design and section 5 presents our experimental results as a number of different findings. Finally, section 6 concludes with a summary and some suggestions for future research.

## 2. Literature

To date, the experimental literature on bank runs has primarily focused on the behavior of depositors in a *single* bank following the set-up of Diamond and Dybvig (1983).[1] These experimental papers have typically focused on the coordination game aspect of that model, asking subjects whether they wish to keep their deposits in the single bank or to withdraw those funds. As in Diamond and Dybvig's model, early withdrawal can be a (self-fulfilling) best response if depositors believe that a sufficient number of other depositors will withdraw early. In particular, Madiés (2006) investigates the possibility and the degree of persistence of self-fulfilling banking panics and shows that those phenomena are both persistent and difficult to prevent. When looking at alternative ways to prevent those type of crises, Madiés's results suggest that a suspension of payment (i.e., more time to think before making a withdrawal decision) is more efficient than partial deposit insurance. Garratt and Keister (2009) show that the frequency of bank runs increases with (1) the uncertainty about the aggregate liquidity demand, and (2) the number of opportunities subjects have to withdraw. Schotter and Yorulmazer (2009) demonstrate that bank runs can be mitigated by the presence of insiders (i.e., depositors who have no uncertainty about the quality of the bank). Arifovic et al. (2013) show that the occurrence of bank runs depends on a coordination parameter, which measures the fraction of depositors that are required to wait so that they can earn a higher payoff than those who withdraw. Their results point towards the existence of three different zones, where bank runs are (i) rare

[1] See Dufwenberg (2015) for a recent survey of the literature.

when the parameter is low; (ii) frequent when the parameter is high; (iii) indeterminate and would depend on the history of the game.

Two papers have used a 2-bank model to explore contagion issues, Brown et al. (2012) and Chakravarty et al. (2015). In both studies, depositors in one bank make their decisions first and the depositors in the second bank observe the decisions of depositors in the first bank before acting. Moreover, the banks' liquidity needs are either linked or independent and this is common knowledge to all depositors in Brown et al. (2012), whereas it is only known to the depositors of the first bank in Chakravarty et al. (2015).[2] While in Brown et al. (2012) contagion occurs only when the banks have economic linkages, in Chakravarty et al. (2015) the depositors' actions in the first bank significantly affect the behavior of depositors in the second bank even when the banks' liquidities are independent. However, with just two banks, network structure cannot play much of a role. None of these papers consider variations in the interbank network structure for efficient risk sharing and the susceptibility of the banking system to financial crises which is the main contribution of this paper.[3] Indeed, Chakravarty et al. (2015, p. 50) conclude by suggesting that for future experimental research on financial crises "there is value not only in reinforcing banking inter-linkages for their value in diversifying risk (as in Allen and Gale, 2000) but also in making those linkages common knowledge". This is precisely the approach that we take in this paper.

## 3. Model and Hypotheses

### 3.1 The Environment

The model is based on the intertemporal model of Allen and Gale (2000). There are three periods, $t = 0, 1, 2$ and four regions. Each of the four regions is served by a local regional bank labeled A, B, C and D. Each region/bank has a continuum of ex-ante identical depositors who have an endowment of one unit of the consumption good at date 0 and nothing for the other two dates. These depositors have preferences as in Diamond and Dybvig (1983); they get utility from consumption only in period 1 (2) with probability w (1-w).

$$U(c_1, c_2) = \begin{cases} u(c_1) & \text{with probability } \quad w \\ u(c_2) & \text{with probability } 1 - w \end{cases}$$

Each bank can invest the deposits of its customers in one of two assets. The liquid (or short) asset acts as storage technology. For each unit of deposits invested in the liquid asset at date $t$, this short-term asset yields a return of 1 at date $t + 1$. The second, illiquid (or long) asset takes two periods to mature, but yields a higher payoff of $R > 1$ per unit invested; if investments in this second asset have to be liquidated early, i.e., in period 1 rather than in period 2, the liquidation return per unit of the asset is given by $r$, where $0 < r < 1$.

[2] Other differences between the papers include the number of withdrawers in each bank, strategic uncertainty regarding types, and the number of repetitions of the game.

[3] Corbae and Duffy (2008) study the role of network structure for equilibrium selection in N-player Stag Hunt games but their main focus is on the endogenous choice of network structure. By contrast, in this paper we impose the network structures exogenously and ask whether those different structures matter for efficiency.

The regions differ in the likelihood that consumers are impatient (early) withdrawers or patient (late) withdrawers. Let $w_i$ denote the probability of early withdrawers in region $i$, and assume that there are just two possible values (low, $w_L$ and high, $w_H$) for this probability: $0 < w_L < w_H < 1$. Assume further that there are two equally likely states of the world, $S_1$ and $S_2$, and that the realizations of the liquidity shocks across the four regions and two states are common knowledge and as given in Table 1:

| | **A** | **B** | **C** | **D** |
|---|---|---|---|---|
| $\boldsymbol{S_1}$ | $w_H$ | $w_L$ | $w_H$ | $w_L$ |
| $\boldsymbol{S_2}$ | $w_L$ | $w_H$ | $w_L$ | $w_H$ |

**Table 1: Distribution of liquidity shocks across banks and states**

The timing of moves is as follows. At date 0, all depositors deposit their endowment in their regional bank and banks invest in the two assets. At date 1, state and depositor uncertainty is resolved; the state of the world is revealed and each depositor's liquidity type is made known according to the probabilities given in Table 1. However, the banks are not able to observe a depositor's type, so it is possible that late (patient) depositors mimic early (impatient) depositors by withdrawing their deposits early.

### 3.2 The Optimal Risk Sharing Contract

The contract that banks in each region offer their depositors at date 0 can be characterized as the solution to a planner's problem that implements the efficient (first best) solution without the need to verify depositor's types. This solution is achievable only if the planner is able to transfer resources across banks in different states of the world. Allen and Gale (2000) further demonstrate that this optimal solution can also be decentralized by the banks themselves through their use of the inter-bank deposit market to insure against uncertain liquidity needs in their own bank/region.

Specifically, in this decentralized setting (which is the environment we study in the laboratory), the optimal contract that each bank offers to its depositors pays $c_1 = 1$ units of consumption to those withdrawing in period 1 and $c_2 = R$ units of consumption to those withdrawing in period 2. Each bank knows that the aggregate demand for liquidity is the same in each state and they also know the average fraction of impatient depositors across all four banks, $\gamma = \frac{w_H+w_L}{2}$. Thus, each bank invests a fraction $\gamma$ of deposits in the short asset and a fraction $1-\gamma$ of deposits in the long asset at date 0, as it is efficient to pay early withdrawers with the small asset and late withdrawers with the long asset.

The optimal risk sharing arrangement is implemented by transfers of resources that banks hold at other banks. For example, if the state of the world $S_1$ occurs, then each bank has $\gamma c_1$ units of the short asset and needs to pay $c_1$ to each impatient depositor at $t = 1$. Banks A and C have excess demand for the short asset in the amount $(w_H - \gamma)c_1$ , while Banks B and D have an excess supply of the short asset in the amount $(\gamma - w_L)c_1 = (w_H - \gamma)c_1$. Thus at date 1, it is possible to satisfy the excess demand of banks A and C if banks B and D transfer their excess supply of the short asset returns, while at date 2, the opposite transfer flow has to take place from banks A and C to banks B and D to satisfy the excess demand in that period. These transfers can be implemented by an appropriate allocation of interbank deposits across the four regions. However the precise amount of these interbank deposits will depend on the network structure of the banking system.

### 3.3 The Importance of Network Structure

With four banks, there are four possible symmetric network structures which are illustrated in Figure 1. While there are also asymmetric, (e.g. "star") network structures, we focus here on *symmetric* network configurations for the banking system as these are easier to explain to experimental subjects and these structures do not involve payoff asymmetries that may trigger inequity (fairness) concerns. The four symmetric network structures examined by Allen and Gale (2000) are as given in Figure 1. Note that for the marriage, incomplete and complete network structures shown in Figure 1 the distribution of shocks given in Table 1 imply that for any state of the world, at least one bank with high liquidity is always connected to at least one bank with low liquidity needs. We will focus on the last two of these symmetric networks in our experiment, the "incomplete" and the "complete" banking system network structures as these are the only two structures that involve network connections among all banks in the economy.

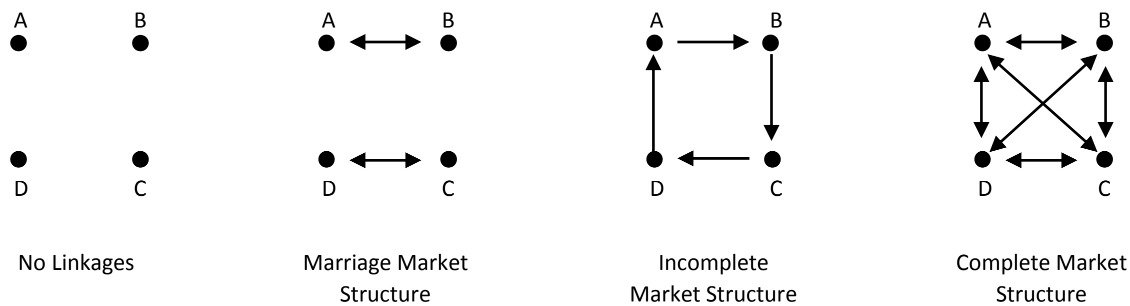


**Figure 1: Symmetric Banking System Network Structures**

In the incomplete network, bank A can place deposits with bank B; bank B can place deposits with bank C; bank C can place deposits with bank D; and bank D can place deposits with Bank A. In line with our experimental instructions, a bank is said to be connected to another bank when it has placed a deposit in that bank. In this incomplete network configuration, Allen and Gale (2000) show that the first best can be achieved if each bank places $(w_H - \gamma)$ deposits in the bank with which they are connected with.

By contrast, under the complete network structure, each bank can place deposits with any of the other three banks. In this case, given the liquidity shock structure of Table 1, each bank's liquidity needs are negatively correlated with two other banks (i.e., a low liquidity bank is correlated with two high liquidity banks and one low liquidity bank, and a high liquidity bank is correlated with two low liquidity banks and one high liquidity bank). It follows that the first best solution can be implemented by having each bank place $(w_H - \gamma)/2$ deposits in each of these other three banks[4].

The main difference between these two networks structures is their susceptibility to what Allen and Gale term a "zero probability at date 0" perturbation. Specifically, suppose there is a state $\bar{S}$ such that the fraction

[4] See Allen and Gale (2000) for a detailed description on how to get these first best allocation.

of impatient depositors in (say) bank/region A is $\gamma + \varepsilon$, so that average liquidity demands across the four regions are higher than in the normal states $S_1$ or $S_2$. As this perturbed state is not known in advance, the continuation equilibrium is different from the normal state and depends on the network structure. Given the zero probability attributed to that state, banks don't change their investment portfolio.

In essence there are three possible outcomes for Bank A. In the first case, it can meet its excess liquidity demands by drawing upon its deposits with other banks and remain solvent. In a second case it can become insolvent, if after withdrawing its deposits from other banks, it must also liquidate some of its position in the long term asset. Finally, a third possibility is that Bank A cannot meet its liquidity needs even by fully liquidating all of its long term asset position and must declare bankruptcy. In the theory of Allen and Gale, the complete network structure is the one that is least susceptible to the last two outcomes, insolvency or bankruptcy, while less connected network structures are more susceptible to these outcomes. Our experiment is designed to test this implication of the theory.

Specifically, we have the following two testable hypotheses (which depend on the parameterization of the model).

First, if the interbank market is *incomplete*, the bank facing the liquidity shock (Bank A) will go bankrupt and the crisis will spread to the interconnected banks. These interconnected banks can also become bankrupt if the liquidation rate, $r$, is small enough. In a similar manner, the crisis then spreads to the whole system. So, if the interbank market is incomplete, and r is sufficiently low, a financial fragility to one bank spreads by contagion to all other banks and leads to an economy wide crisis.

Second, if the interbank market is *complete*, the initial impact of a financial shock in one bank may be mitigated if every bank takes a small hit (that is, every bank liquidates some of the long asset). This possibility exists regardless of the value of the liquidation rate, $0 < r < 1$. Of course it is also possible that depositors refuse to accept such losses (withdraw early) and a contagious wave of bankruptcies occurs in this setting as well. However, the possibility that the crisis is localized to the shocked bank (A) is an equilibrium possibility and that is the main difference that we wish to test with our experimental design.

### 3.4 Payoffs

Following Allen and Gale (2000), impatient depositors do not earn any interest on their deposits and are therefore promised a return exactly equal to the amount initially deposited at their bank. On the other hand, patient depositors are promised a return $R$. However, due to the financial perturbation and possible contagion, some or all banks will need to liquidate part or all of their long-term asset. Given that this is done at a liquidation cost of $0 < r < 1$, those banks will be unable to pay back, either patient, impatient, or both type of depositors, their promised returns. In those situations, the payoffs will be given by the ratio of the total asset value of the bank and the number of withdrawers (including both depositors and connected banks).

Let $q^i$ represent the value of all deposits in bank $i$ at $t = 1$. If $q^i$ is less than the promised return in period 1 (i.e. $c_1$), then whether the withdrawer is another bank or a consumer, they will each get $q^i$ from the bank for each unit invested at $t = 0$. Given the nature of the interbank network structures, this requires all $q^i$ values to be determined simultaneously.

Consider for instance the *Incomplete Market Structure*. Assume that all depositors of Bank A withdraw at $t = 1$, so that the total demand is $n + z$, where $n$ is the number of withdrawers and each deposited one unit at $t = 0$ and $z$ is the deposit that Bank D holds in Bank A. The liabilities of Bank A are then valued at:

$$(n + z)q^A \tag{1}$$

The assets of this bank are the *y* units of short asset, the *x* units of the long asset that will be prematurely liquidated at rate, *r*, and the *z* deposit in Bank B. Therefore, the assets value is given by:

$$y + rx + zq^B \tag{2}$$

The equilibrium values of $q^A$ require the values of assets and liabilities to be equal, so that:

$$q^A = \frac{y+rx+zq^B}{n+z} \tag{3}$$

A similar equation will hold for any Bank *i* in which $q^i$ is less than the promised return in period 1.

This equation can be used as long as $q^B$ is equal to the promised return, $c_1$. If $q^B < c_1$, then the equivalent equation is needed to estimate $q^B$, which will depend on the value of $q^C$; and so on.

Similarly, for the *Complete Market Structure*, if all depositors of Bank A withdraw at $t = 1$, so that the total demand is $n + 3z$, where $n$ is the number of withdrawers and each deposited one unit at $t = 0$ and $z$ is the deposit that Banks B, C, and D hold in Bank A. The liabilities of Bank A are then valued at:

$$(n + 3z)q^A \tag{4}$$

The assets of this bank are the *y* units of short asset, the *x* units of the long asset that will be prematurely liquidated at rate, *r*, and the *z* deposit in each of Bank B, C, and D. Therefore, the assets value is given by:

$$y + rx + z(q^B + q^C + q^D) \tag{5}$$

The equilibrium values of $q^A$ require the values of assets and liabilities to be equal, so that:

$$q^A = \frac{y+rx+z(q^B+q^C+q^D)}{n+3z} \tag{6}$$

A similar equation will hold for any Bank *i* in which $q^i$ is less than the promised return in period 1.

This equation can be used as long as $q^B$, $q^C$, and $q^D$ are equal to the promised return, $c_1$. If $q^B < c_1$, and/or $q^C < c_1$, and/or $q^D < c_1$, then the equivalent equation is needed to estimate $q^B$ and/or $q^C$ and/or $q^D$, which will depend on the value of all deposits at $t = 1$ in the banks they are connected with.

## 4. Experimental Design and Procedures

In our experimental setting, as in the theory, there are four banks labelled A, B, C and D. Each participant in our experiment is assigned the role of a depositor in one of these four banks. The experimental setting and payoffs have been determined based on 4 depositors for each bank.

We set the probability that a depositor is impatient in a bank that faces a low or high liquidity shock to $w_L = 1/4$ and $w_H = 3/4$, respectively. Thus, the average fraction of impatient depositors in the economy is $\gamma = \frac{\frac{1}{4}+\frac{1}{4}+\frac{3}{4}+\frac{3}{4}}{4} = \frac{1}{2} = 0.5$. Therefore in the perturbed state $\bar{S}$, among the banks that do not face the liquidity shock, the number of impatient depositors is equal to $4 * \gamma = 2$, while in the bank that faces the liquidity shock it is $4 * (\gamma + \varepsilon) = 3$.

Knowing the average fraction of impatient withdrawers, the banks will invest a fraction $\gamma = 0.5$ of deposits in the short asset and a fraction $1 - \gamma = 0.5$ in the long asset. As previously mentioned, Allen and Gale (2000) demonstrate that the first best is achieved by placing a fraction $(w_H - \gamma) = \frac{3}{4} - \frac{1}{2} = \frac{1}{4}$ and $(w_H - \gamma)/2 = \frac{1}{8}$ of the deposits in the bank(s) with which they are connected in the incomplete and complete market structures, respectively. Moreover, in this experiment we set $R = 2$ and $r = 0.2$.[5] All those parameters enable us to then simultaneously solve for $q^i$, the value of all deposits in bank $i$ at $t = 1$, to estimate the payoffs offered to depositors when banks are illiquid or bankrupt (i.e., they are unable to pay the promised return).

While the experimental setting and payoffs are established on the basis of 4 depositors in each bank or 16 depositors for the total economy, our primary focus is the behaviour of ‘strategic’ players, namely, the patient depositors who can choose whether or not to withdraw early, that is, in period 1; the impatient depositors just mechanically withdraw early and so are of little behavioural interest. Therefore, we parameterized our experiment in such a way that human subjects are only needed to play the role of patient depositors in all banks. That is, in our experiment, each economy consists of just 8 human subject depositors, 2 for each bank. Effectively, the actions that would be taken by the early withdrawers are built into the payoff structure. In addition, one of these “patient” human depositors was randomly subjected to the perturbed shock and forced to withdraw early in period 1. To further simplify the experimental setting, the financial fragility shock would always originate within bank A.

An experimental observation consists of the play of 30 rounds by the same 8 subjects representing a single economy. At the beginning of each round, each pair of subjects is randomly assigned to 1 of the 4 banks. We chose a fixed matching design to better allow for learning behavior, while the random assignment of players to banks at the start of each round was chosen to avoid having the same subjects be repeatedly exposed to the liquidity shock in Bank A.

Specifically, in period $t = 0$, each group is assigned to each bank and is informed about which bank it has been assigned to. We used the terminology “group” to refer to the 2 subjects assigned to each of the banks. Next, subjects deposit their endowment of 100 experimental pounds (EP) in their bank.[6] Then, in period $t = 1$, depositors in bank A learn whether or not they are the one forced to withdraw their deposit. All other 7 patient depositors have to make a single decision: whether to withdraw their deposit in $t = 1$ or wait until $t = 2$ (not withdraw). Participants have full information about the perturbation shock, that is, that the shocked bank is always bank A and to which bank they are assigned in each round. After all decisions are

[5] The numbers are chosen so as to test Allen and Gale’s hypothesis that for a high pre-mature liquidation cost, the incomplete market structure could face an economy wide bankruptcy, whereas this effect is mitigated with a complete market structure where risk sharing takes place.
[6] This was computerized and the subjects’ endowments were automatically deposited in their respective bank.

made (i.e., withdraw or not withdraw), each participant learns the outcome of their decision (round earnings), the decision of the group-mate and the aggregated decisions of the depositors in the connected banks, a history table summarizing this information for previous rounds is also presented. Round earnings are in experimental pounds. At the end of the experiment, one round is randomly selected for payment.

Round earnings depend on the network structure. We implement each network structure as a separate between-subjects treatment. In each treatment subjects are informed of the network linkages and the payoff consequences from choices by their own bank members and others via the interbank network connections. In our setup, there are 54 possible combinations of withdrawal requests. To ease presentation, the payoff tables were created by grouping the combinations according to the choice of the person in the same bank (group-mate) and the number of withdrawal requests in the connected banks and taking the median of the payoffs for each grouping, as shown on Tables 2 and 3 below.[7]

For the incomplete network structure, the payoff table for patient depositors was as shown in Table 2:

| | *Choice of your group-mate* | **N** | | | **W** | | |
|---|---|---|---|---|---|---|---|
| | *Number of Withdrawals in the connected bank* | **N** | **1W** | **2W** | **N** | **1W** | **2W** |
| *Your Choice* | **N** | 100 | -4 | -83 | -100 | -100 | -100 |
| | **W** | -15 | -20 | -24 | -32 | -36 | -39 |

**Table 2: Incomplete Network Payoff Table**

The player's own choice, to Not Withdraw (N) or to Withdraw (W) is indicated in the left column and the choice of the group-mate (in the same bank) is indicated in the top row of the two right columns. These two right columns are further sub-divided up according to the choices made by the two depositors in the connected bank - the other bank holding deposits of the bank the two players are in. Here, N means no (0) Withdraw choices by depositors in the connected bank, while 1W and 2W mean 1 or 2 withdrawals, respectively by the depositors in the connected bank. Recall that initially all subjects deposit 100 EP in the bank, so the payoff numbers indicate the *additional* EP from the various actions. Thus, if the group-mate does not withdraw, N, nor are there any withdrawals in the connected bank, N (a possibility so long as the connected bank is not Bank A), then, the players would gain 100 EP from choosing N (not withdraw) reflecting our choice of R = 2. In this same scenario, the player would lose 15 EP if she instead chose to withdraw W for a net payoff of 85 EP. Notice that losses are always capped at -100, resulting in a net payoff of 0. Payoffs less than 100 reflect the liquidation rate choice of r = 0.2. Finally, notice that if a player's group-mate withdraws, it is always a dominant strategy to withdraw as well. In Bank A, one depositor is forced to withdraw and so his patient-type group mate, knowing that he is in Bank A and facing the payoff table above should play a best response of withdrawing as well. Since there will (rationally) be two withdrawals (2W) in Bank A, members of the bank connected to (with deposits in) Bank A, i.e. Bank D depositors, in the incomplete network structure, should rationally anticipate that they will face 2W in the connected bank, in which case the dominant strategy is for both players in Bank D to choose W.

[7] The payoff tables for all possible combinations of withdrawal demands are available upon request.

Recognizing that the two players in Bank D will play 2W, the two players in Bank C should also play W, and, recognizing this outcome, the two players in Bank B will also play 2W, making the financial contagion complete. Thus for the incomplete market structure we have the following hypothesis:

***Hypothesis 1:*** *In the INCOMPLETE market structure, the original financial shock spreads to all banks as one after the other face bankruptcy.* [8]

For the complete market structure, the payoff table for patient depositors was as shown in Table 3:

| | *Choice of your group-mate* | **N** | | | | | | | **W** | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | *Number of withdrawals in the connected banks* | **0W** | **1W** | **2W** | **3W** | **4W** | **5W** | **6W** | **0W** | **1W** | **2W** | **3W** | **4W** | **5W** | **6W** |
| *Your Choice* | **N** | 100 | 67 | 25 | -18 | -77 | -100 | -100 | -100 | -100 | -100 | -100 | -100 | -100 | -100 |
| | **W** | -13 | -15 | -17 | -19 | -22 | -24 | -26 | -29 | -31 | -33 | -34 | -37 | -38 | -40 |

**Table 3: Complete Network Payoff Table**

The payoff table for the complete network case is read similarly to the payoff table for the incomplete network case, but now the possible actions of players in the connected banks is greater, as there are three connected banks in the complete network and so anywhere from 0 to 6 players can choose to withdraw among these connected banks. Of course 0W by members of connected banks is only a possible outcome for members of Bank A, and this fact was made clear in the instructions.

Notice under the complete network structure that, conditional on one's group-mate not withdrawing (i.e. choosing N), it is a best response for the player not to withdraw so long as the number of withdrawals by players in connected banks does not exceed 3W. Since for banks B-D, the number of withdrawals in connected banks can rationally be expected to be, at minimum, 2 - namely the two players in Bank A - the efficient equilibrium is for no members of Banks B-D to choose withdraw (W) and as a result all players in Banks B-D earn 25 EP on top of their 100 EP investment. On the other hand, if players believe that all others will choose to withdraw, then it is an equilibrium best response for all players to withdraw as well, so that the inefficient financial contagion equilibrium also exists in this complete network setting as well. Assuming that payoff efficiency is the relevant equilibrium selection criterion, we have the following hypothesis:

[8] Given our experimental setup, a bank will face bankruptcy if the number of withdrawers in that bank is higher than the average fraction of impatient depositors across all four banks, $\gamma$. This implies that Bank A which faces the financial shock is always bankrupt, while it takes one patient depositor in any other bank to imitate the behavior of impatient depositors for that bank to become bankrupt.

***Hypothesis 2:*** *In the COMPLETE network, only the bank facing the financial shock will go bankrupt. The financial crisis does not become global in a fully integrated financial system.*

The experiment was computerized using zTree (Fischbacher, 2007). Subjects were students at the University of East Anglia. No subject had any prior experience with our experimental design and subjects were only allowed to participate in a single session/treatment of our experiment. We obtained 11 observations on choices by 8-subject cohorts under the incomplete network treatment and 11 observations on 8-subject cohorts under the complete network treatment. Thus our study involved a total of 176 subjects (average age = 21.9 years; 49.5% females).

Each experimental session began with subjects being given written instructions which were then read aloud in an effort to make those instructions common knowledge (copies of the instructions are included in the Appendix). After the instructions were read, subjects had to answer a number of questions designed to check their comprehension of the written instructions. Subjects who made mistakes were instructed as to the correct answers prior to the first round of the game.

After completing 30 rounds, subjects were informed which round was randomly selected for payment and the corresponding payoff for that round. Earnings in experimental points (EP) were converted into British Pounds at an exchange rate of 1 EP = 0.1 British Pounds. In addition, subjects received a 3 British pounds show-up fee. Between the end of the experiment and the payment phase, a demographics and feedback questionnaire was administered. Participants received their payment in cash at the end of the session.

Each session was completed within 1 hour and 2 – 3 cohorts (16 – 24 subjects) participated in a session. The average payment was 8.96 British pounds (SD 2.96 British pounds) in the incomplete network treatment and 9.65 British pounds (SD 4.47 British pounds) in the complete network treatment.

## 5. Findings

We report the results of our experiment as a number of different findings. We begin with an analysis of the main treatment difference between the incomplete and complete network structures before moving on to more micro level differences.

***Finding 1:*** *Consistent with Hypothesis 1, in the INCOMPLETE network structure, the original financial shock frequently spreads to the other three banks, which then face bankruptcy.*

Support for Finding 1 comes from Figure 2 which shows the number of banks other than Bank A experiencing bankruptcy in each of the 30 rounds in the incomplete network treatment. Here bankruptcy is defined as an inability to meet the payment promises made to depositors choosing to withdraw at $t = 1$.

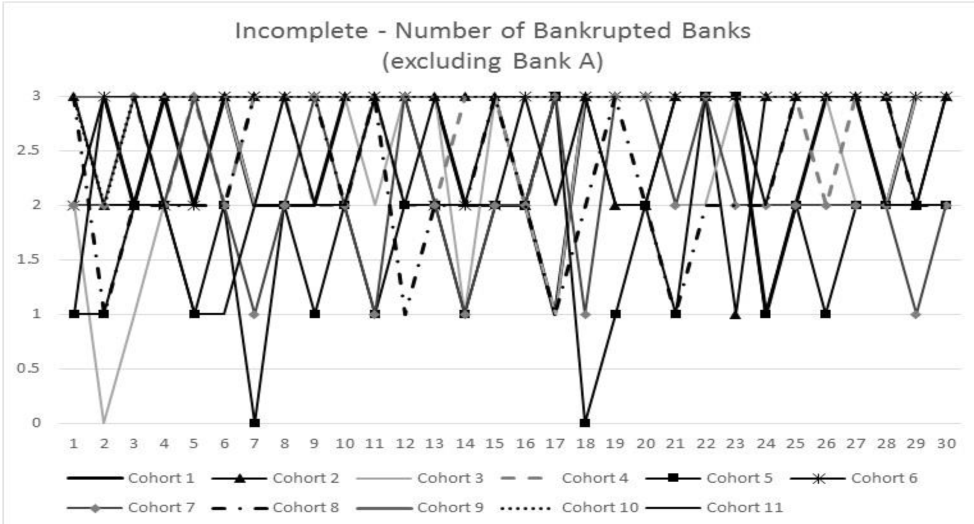


**Figure 2: No. of Banks B-D Bankrupted Each Round, All 11 Cohorts of the Incomplete Treatment**

Figure 2 reveals that the number of bankrupted banks was typically between 2 and 3 (the mean is about 2.5), so that the financial contagion was not always perfectly complete. Nevertheless, it is clear from the figure that the outcome where 0 of the banks other than Bank A were bankrupt was a rare occurrence, and never occurred in the final 10 rounds of play.

***Finding 2:*** *Inconsistent with Hypothesis 2, under the COMPLETE network, the financial crisis sometimes becomes global, spreading to all 3 banks and sometimes it is contained so that only the bank facing the financial shock goes bankrupt.*

Support for Finding 2 comes from Figure 3 which shows the number of banks other than Bank A experiencing bankruptcy in each of the 30 rounds in the complete network treatment.

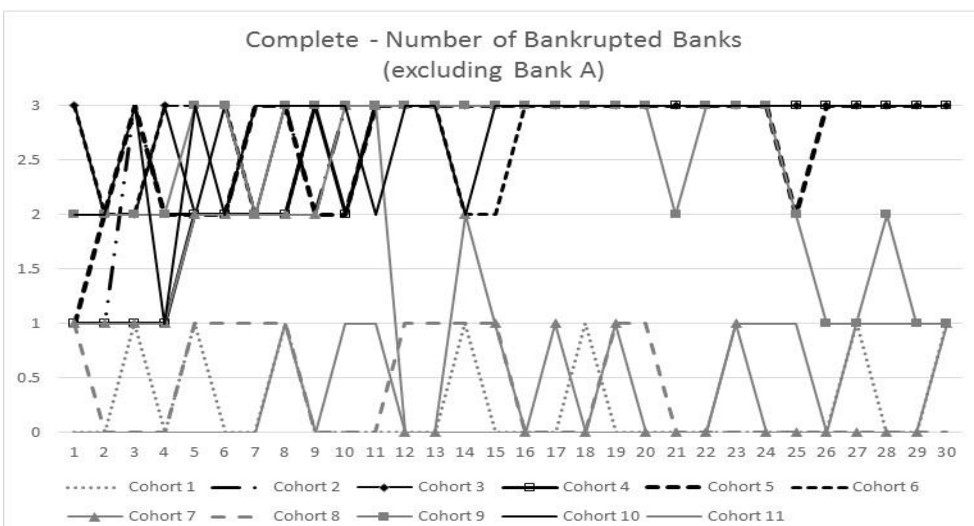


**Figure 3: No. of Banks B-D Bankrupted Each Round, All 11 Cohorts of the Complete Treatment**

Figure 3 clearly reveals a bifurcated outcome, with 5 out of 11 cohorts repeatedly experiencing 0 or 1 bankruptcies among banks other than Bank A—a “contained” financial crisis—while the other 6 cohorts often experience a complete, or nearly complete global financial contagion where all or nearly all banks (other than Bank A) immediately become bankrupt. While both pure equilibria are possibilities under the complete network structure, our results indicate that complete interbank connectedness and the more efficient risk sharing that it allows for, provides *no guarantee* that agents will coordinate on the efficient (first best) outcome.

Figure 4 provides a further comparison of the impact of our two network treatments on bankruptcies beyond Bank A. This figure reports the mean number of bankrupt banks (other than Bank A) from all cohorts in the incomplete network treatment along with the 95% confidence intervals. For the complete network treatment, the figure reports the mean number of bankrupt banks for the 5 cohorts that succeeded in largely containing the crisis to Bank A (“good” equilibrium cohorts) and the mean number of bankrupt banks for the 6 cohorts that did not succeed in containing the crisis (“bad” equilibrium cohorts), again along with the 95% confidence intervals. Figure 4 reveals that over time, there is clear separation in the number of bankrupt banks between the good and the bad cohorts of the complete network treatment. The figure further reveals that there is no significant difference in the number of bankrupt banks between the “bad” cohorts of the complete network treatment and all 11 cohorts of the incomplete network treatment.

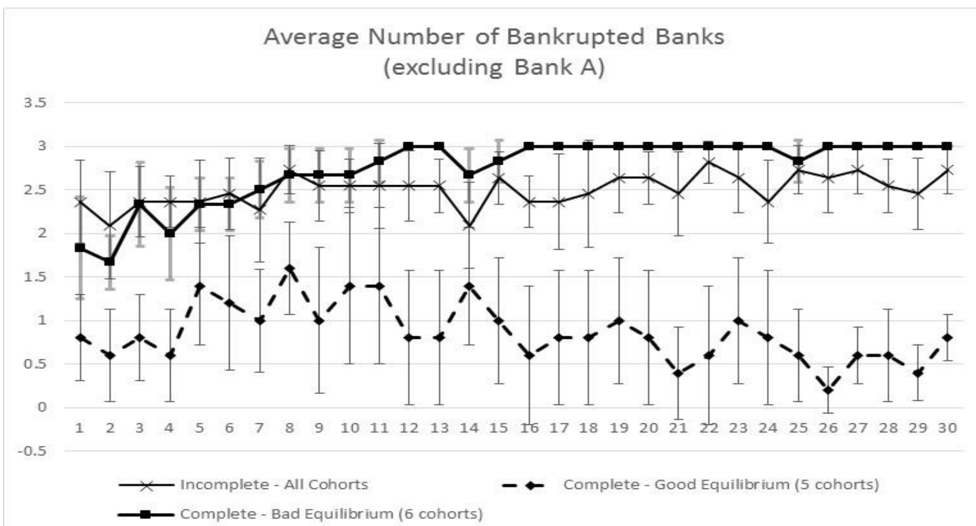


**Figure 4: Mean and Variance in the Number of Bankrupted Banks over Time, Incomplete and Complete Network Treatments**

Figures 5abcd shed some light on the mean *withdrawal* decisions by cohorts at banks, A, B, C, and D over the 30 rounds of each session using 5-round moving averages on the number of withdrawals per bank.


**Figures 5abcd: Withdrawal decisions by Banks in Various Treatments/Cohorts over Time**

The top two panels, Figures 5ab, clearly reveal that under the incomplete network structure, the number of withdrawal requests in Bank A is the greatest on average, but closely followed by high (greater than 1 on average) withdrawal requests in Banks D, C and B. By contrast under the complete network structure, there is much clearer separation in mean withdrawal requests; the two withdrawal outcome, associated with bankruptcy, is clearly greater for Bank A than for the other three banks suggesting that under the complete network structure, the contagion is contained to some extent.

The bottom two panels, Figures 5cd, again make a distinction between the 5 cohorts of the complete network treatment that coordinated on the good risk sharing outcome and the 6 other cohorts that did not. The separation in the mean withdrawal requests between Bank A and Banks B, C, and D is clearly evident in Figure 5c (bottom left) among the cohorts of the complete network treatment for which the financial contagion was contained to Bank A. For the other 6 cohorts of the complete network treatment, Figure 5d (bottom right) indicates very little difference over time in the mean number of withdrawal requests, which are all close to 2 by the final 30th period (indicating perfect bankruptcy). These figures provide additional support for Finding 2.

Most importantly, the 5-round moving average enables us to smooth out the effect of the withdrawal requests for each round. A clear difference emerges between the incomplete and the complete network structures. For the latter, Banks B, C, and D withdrawal requests seem to overall be synchronized, leading to a global financial crisis. However, for the former a spillover effect is observed: the financial shock of Bank A is transmitted to Bank D (which has invested its deposit into Bank A). This then spreads to Bank C, who has invested in Bank D; and finally, to Bank B which has invested in Bank C. There is therefore a clear contagion in the form of spillover in the incomplete market structure, before the crisis becomes global.

Given Findings 1-2, it immediately follow that efficiency can be greater under the complete than under the incomplete network structures which we summarize as follows:

*Finding 3: Efficiency differences between the two treatments depend on whether cohorts in the Complete network structure treatment succeeded in containing the contagion or not.*

Support for Finding 3 comes from Figure 6 which shows average efficiency and the 95% confidence intervals for each of the 30 rounds of the experiment using data from 1) all 11 cohorts of the incomplete network treatment, 2) the 5 cohorts of the complete network treatment that were able to contain the financial contagion and 3) the 6 cohorts of the complete network structure for which the contagion was global.

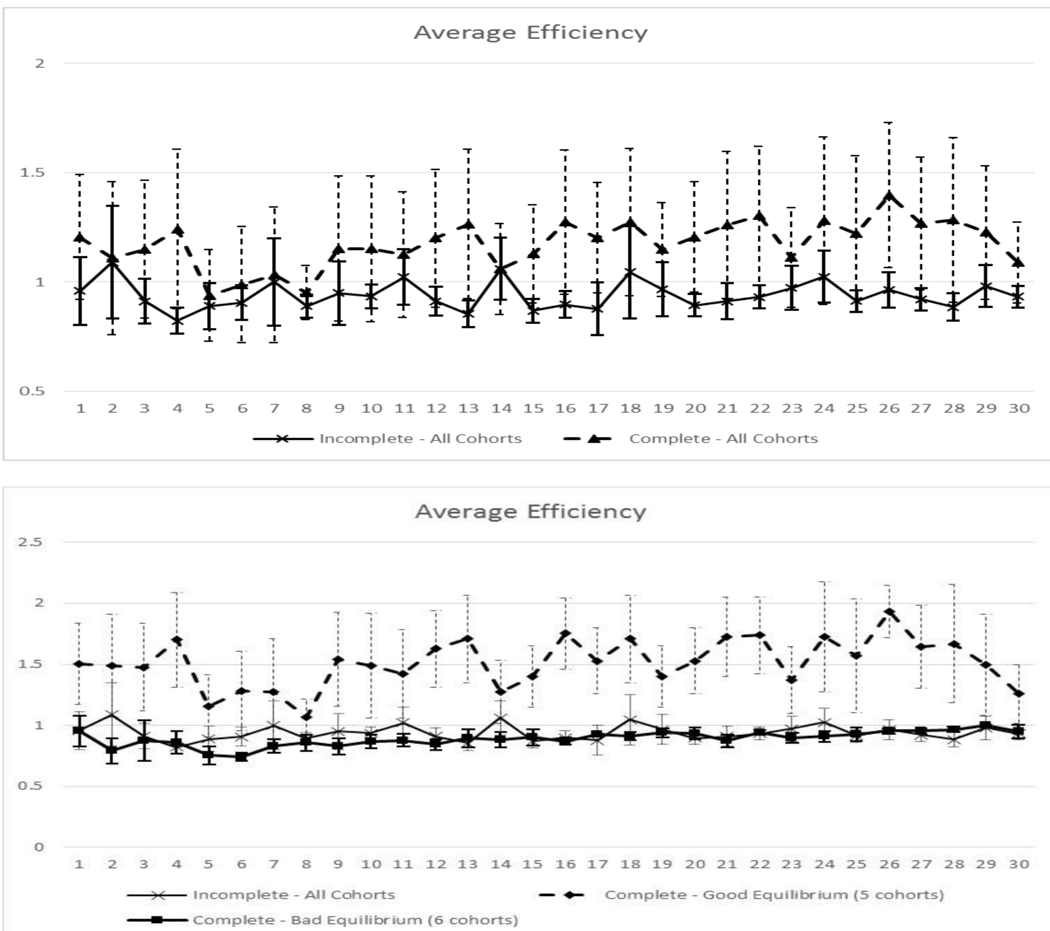


**Figure 6: Average Efficiency over Time in the Incomplete and Complete Treatments. The Complete Treatment is subdivided into Cohorts closer to the Good or to the Bad Equilibria of the Model.**

We only find efficiency differences between our two treatments if we make this distinction among cohorts in the Complete Network treatment. The Figure reveals that the 5 cohorts that achieved outcomes approximating the efficient (Good) risk sharing equilibrium have average efficiency levels that are significantly greater than all cohorts in the Incomplete Network treatment and those cohorts in the Complete Network treatment that did not succeed in containing the financial contagion.

We next turn to an analysis of *individual* withdrawal decisions using a mixed effects panel logit regression estimator with 3 levels: Individual, Group and Cohort. For this exercise we consider all individual withdrawal decisions, where the subject could choose whether or not to withdraw. Thus we exclude the withdrawal decisions of those subjects who were assigned to Bank A and were the one member of Bank A who was *forced* to withdraw; there is 1 such subject in each 8-subject cohort who fits this description in each round. Thus, we have data on the voluntary withdrawal decisions of 7 subjects per cohort over 30 rounds and we have 22 cohorts in total (11 of each treatment). This provides us with $\mathbf{7 \times 30 \times 22 = 4,620}$ observations on individual withdrawal decisions from our experiment.

The results of our regression exercise are reported in Table 4, where the dependent variable is always the individual withdrawal decision in each round, 0 = not withdraw (wait until period 2) and 1 = withdraw early (in period 1). Thus we adopt a mixed effect logit estimator and Table 4 reports marginal effects. The explanatory variables are as follows: Incomplete is a dummy variable for whether choices were made under the incomplete network treatment, Withdraw in $\boldsymbol{t-1}$ is the lagged withdrawal decision (if the player is not the one in Bank A who was forced to withdraw in the previous round), Partner withdraw in $\boldsymbol{t-1}$ is the decision of the player's partner in the lagged round, Number of withdraws in connected banks in $\boldsymbol{t-1}$ is the normalized number of withdrawals in connected banks in the lagged round.[9] In addition to these economic choice variables, we also include a number of control variables making use of demographic data we collected on individual subject characteristics. These include sex, age, English language skills and prior experience in economic decision-making (DM) experiments. We find that most of latter demographic factors have no significant explanatory power on withdrawal decisions except for age.

[9] Recall that in the complete network for any bank, the number of connected banks is three while in the incomplete network is one, therefore in the complete network the maximum number of withdrawals in the connected banks is six whereas in the incomplete network is two. To make treatments comparable, we used a unity-based normalization, that is: $x' = \frac{x-\min(x)}{\max(x)-\min(x)}$, where $x$ is the observed number of withdrawals in the connected banks.

| VARIABLES | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
|---|---|---|---|---|---|---|---|---|
| If treatment = Incomplete | 1.258** | 1.230** | 1.147** | 1.129** | 1.058** | 1.037** | 0.882** | 0.858** |
| | (0.573) | (0.568) | (0.513) | (0.508) | (0.473) | (0.468) | (0.388) | (0.383) |
| Withdraw (if not forced) in t-1 | | | 0.731*** | 0.731*** | 0.561*** | 0.561*** | 0.504*** | 0.505*** |
| | | | (0.0982) | (0.0982) | (0.102) | (0.102) | (0.103) | (0.103) |
| Partner withdraw in t-1 | | | | | 0.762*** | 0.757*** | 0.771*** | 0.768*** |
| | | | | | (0.104) | (0.104) | (0.105) | (0.105) |
| Number of withdrawals in connected banks in t-1 (normalized) | | | | | | | 1.199*** | 1.200*** |
| | | | | | | | (0.168) | (0.168) |
| if sex = Female | | 0.139 | | 0.142 | | 0.153 | | 0.142 |
| | | (0.186) | | (0.165) | | (0.183) | | (0.189) |
| Age (years) | | -0.0412** | | -0.0311* | | -0.0341* | | -0.0356* |
| | | (0.0194) | | (0.0167) | | (0.0190) | | (0.0195) |
| If native language not English | | 0.127 | | 0.181 | | 0.152 | | 0.173 |
| | | (0.215) | | (0.190) | | (0.210) | | (0.216) |
| Experience in DM Experiments | | -0.00631 | | -0.00660 | | -0.00449 | | -0.00327 |
| | | (0.0206) | | (0.0182) | | (0.0203) | | (0.0210) |
| Constant | 0.159 | 0.928 | -0.129 | 0.407 | -0.481 | 0.119 | -1.163*** | -0.537 |
| | (0.404) | (0.576) | (0.364) | (0.506) | (0.339) | (0.525) | (0.296) | (0.508) |
| Observations | 4,620 | 4,620 | 3,904 | 3,904 | 3,904 | 3,904 | 3,904 | 3,904 |
| Number of groups (cohorts) | 22 | 22 | 22 | 22 | 22 | 22 | 22 | 22 |

Notes: Marginal effects from mixed-effects panel logistic regression, 3 levels: Cohort (N=22)-group (N=88)-individual (N=176). Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

**Table 4: Mixed effects Panel Logit Regression Analysis of Withdrawal Decisions**

Table 4 reveals two other interesting results. First, the coefficient on the treatment dummy variable is found to be statistically significant, which indicates that the *incomplete* network structure makes it more likely that subjects will choose to withdraw early (in period 1) relative to the baseline, complete network structure. This result provides further evidence in support of Findings 1 and 2 that the interbank network structure matters for the incidence of contagion.

Second, Table 4 further reveals that *history* also matters, as the lagged withdrawal choices by the subject, his partner in the previous round or by depositors in the connected banks of the previous round all increase the likelihood that the subject chooses to withdraw in the current round. This finding suggests that there is some path dependence of withdrawal outcomes from prior rounds that accounts for the withdrawal decisions of cohort members in the current round, but even accounting for this fact, the network structure still matters for the frequency of current withdrawal choices. We summarize these findings as follows:

*Finding 4: Network structure matters for individual withdrawal decisions even after controlling for demographic factors and the immediate prior history of withdrawal decisions.*

Finally, we consider whether a financial contagion, if it occurs, unfolds in the manner predicted by the theory under the two different network structures. Recall that under the complete network structure, Bank A has interbank connections with the other three banks. Thus, a bankruptcy in Bank A has immediate payoff consequences for depositors in all three of the other banks. If depositors in these other banks do not all immediately choose to withdraw, they can achieve the first best equilibrium wherein the financial shock is localized to Bank A. However, if depositors believe that enough other non-bank-A depositors will withdraw early, the contagion to all depositors withdrawing should occur simultaneously and with the same incidence across all three banks. Under the incomplete network structure, the bankruptcy of Bank A (if it occurs) has immediate spillover effects only to depositors in Bank D, which by design, holds some of its depositors' deposits in Bank A. If Bank D fails, then depositors in Bank C are adversely affected, and if Bank C fails depositors in Bank B are adversely affected completing the financial contagion around the incomplete network. Of course, in a financial contagion this wave of bankruptcies should also play out immediately in period 1, but because of the incompleteness of the network structure it may well be that distance from the source of the financial crisis-- namely Bank A-- matters for the timing of withdrawal decisions.

To examine whether network structure matters for the speed with which a contagion unfolds we examine the decisions of non-Bank-A depositors to ``wait'' (i.e., to not withdraw = 1) under the two different network environments. Specifically, we considered the impact of depositors' distance from bank A on their waiting decision. We used the same mixed effects panel logit estimator as in Table 4 to examine the waiting choices of non-Bank-A depositors as a function of dummy variables, B, C, D, representing their bank membership. The results (marginal effects) are reported in Table 5.

| VARIABLES | Incomplete | Complete |
|---|---|---|
| Bank | | |
| B | 2.349*** | 3.027*** |
| | (0.192) | (0.192) |
| C | 2.006*** | 3.086*** |
| | (0.189) | (0.192) |
| D | 1.522*** | 3.025*** |
| | (0.194) | (0.192) |
| Constant | -3.295*** | -2.961*** |
| | (0.365) | (0.491) |
| Observations | 2640 | 2640 |
| Number of groups (Cohorts) | 11 | 11 |

Notes: marginal effects from mixed-effects panel logistic regression, 3 levels: cohort (N=22)-group (N=88)-individual (N=176). Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

**Table 5: Mixed effects Panel Logit, showing marginal effects for Banks (Baseline Bank A) by treatment Dependant variable "wait=1"**

The table reveals that for the incomplete network treatment, the coefficients (marginal effects) are different between the three banks. Indeed, a Wald test of the hypotheses that estimated coefficients are equal between banks D and C; between banks C and B; and between banks D and B is easily rejected (Prob > $\chi^2$ < 0.05 for all three pairwise comparisons). Notice further that waiting times are *lower* the closer is the connection to Bank A (the baseline). In particular, bank waiting times are such that the order across the three banks is $D < C < B$. This pattern is consistent with a behavioural bias favouring (against) early withdrawal the more (less) directly connected the bank is to the source of the financial crisis, namely Bank A.

Conversely, for the complete network treatment, coefficients are *not* significantly different from one another across the three banks; (Wald Test Prob > $\chi^2$ > 0.10 in all three pairwise comparisons) and are *higher* than in the incomplete network. The lack of a difference in coefficients across the three banks is consistent with the theoretical prediction that a contagion, if it happens, does so instantaneously across all four banks in the economy. The higher marginal effects in the complete network treatment relative to the incomplete network treatment simply reflect the finding that financial contagions always occur under the incomplete network structure, but occur less frequently under the complete network structure. We summarize the results from Table 5 as follows:

*Finding 5: While theory predicts that a financial contagion, if it occurs, spreads to all banks immediately in period 1, in the experiment we find that the contagion is slower to unfold in the incomplete network structure and the timing of depositors' decisions to wait/withdraw depends on the distance of their bank from the source of the financial crisis (Bank A).*

## 6. Conclusions and Extensions For Future Research

Modern banking systems involve many connections across banks (and non-banks) e.g. for risk management and payment processing reasons. Thus, it is not surprising that modern financial crises will potentially have contagion effects with spillovers from one bank to another. In this paper we report on the first experiment exploring the role of interbank network structure for the incidence of financial contagion. Consistent with theoretical framework of Allen and Gale that we implement and test in the laboratory we find that financial contagions are common under incomplete network structures in cases where efficient risk sharing is not possible. We further find that when efficient risk sharing is possible under a complete interbank network structure, financial contagions can sometimes be contained. However, we also find that financial contagions continue to be a possibility even under the complete interbank network structure. Thus, an important implication of our results is that while more complete interbank network structures may reduce the incidence of financial contagions by facilitating more efficient risk sharing among banks, such complete network structures are not a *panacea* for preventing such contagions.

There are several directions for future research on this topic. First, we have only considered a single parameterization of the model's parameters, e.g. R, r, and the number of banks in the banking system. It would be of interest to consider variations in these parameters in combination with the variation in network structure that we do implement in our experiment. In particular, one could choose a higher liquidation rate r than is used in our study, which would allow for the possibility that financial crises could be contained to Bank A even under the incomplete network structure, although a global financial crisis would also remain a possibility under this same parameterization.

Second, one could consider more realistic asymmetric interbank network structures where different banks have differing numbers of interbank network connections that could be determined according to existing, real-world interbank network structures.

Finally, an implication of our last Finding 5 is that there may be some value to modeling financial contagion among interconnected banks in incomplete network environments using an explicitly *dynamic game* approach as opposed to the static, simultaneous-game approach used in our experiment based on the model of Allen and Gale (2000).

We leave these extensions to future research.

## 7. References

Allen, F., and D. Gale (2000), “Financial Contagion,” *Journal of Political Economy* 108, 1-33.

Arifovic, J., J.H. Jiang and Y. Xu (2013), “Experimental Evidence of Bank Runs as Pure Coordination Failures,” *Journal of Economic Dynamics and Control* 37, 2446–2465.

Bock, O., Baetge, I., and Nicklisch, A. (2014). “hroot: Hamburg registration and organization online tool.” *European Economic Review*, 71, 117-120.

Brown, M., Trautmann, S. and Vlahu, R. (2012), “Contagious bank runs: experimental evidence”, Working Paper No. 363, De Nederlandsche Bank (DNB), Amsterdam.

Chakravarty, S., M.A. Fonseca and T.R. Kaplan (2015), “An Experiment on the Causes of Bank Run Contagions.” *European Economic Review*, 72, 39-51.

Corbae, D. and J. Duffy (2008), “Experiments with Network Formation,” *Games and Economic Behavior* 64, 81-120.

Diamond, D.W. and P. Dybvig (1983), “Bank Runs, Deposit Insurance and Liquidity,” *Journal of Political Economy* 91, 401-419.

Dufwenberg M. (2015), “Banking on experiments?” *Journal of Economic Studies* 42:6, 943-971.

Fischbacher, U. (2007), “zTree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics* 10, 171-178.

Garratt, R. and T. Keister (2009), “Bank Runs: An Experimental Study,” *Journal of Economic Behavior and Organization* 71, 300-317.

Gorton, G.B. (2010), *Slapped by the Invisible Hand: The Panic of 2007* Oxford: Oxford University Press.

Kiss, H., I. Rodriguez-Lara and A. Garcia (2012), “On the Effects of Observability and Deposit Insurance on Bank Runs: An Experimental Study,’’ *Journal of Money, Credit and Banking* 44, 1651-1665.

Madiés, P. (2006), “An Experimental Exploration of Self-Fulfilling Banking Panics: Their Occurrence, Persistence and Prevention,” *Journal of Business* 79, 1831-1866.

Schotter, A. and T. Yorulmazer (2009), “On the Dynamics and Severity of Bank Runs: An Experimental Study” *Journal of Financial Intermediation* 18, 217-241.

## Appendices

### A.1 Printed instructions: Incomplete network structure

**Instructions**

Welcome to this experiment in economic decision-making. Please pay careful attention to these instructions as they explain how you earn money from the decisions that you make. After we read the instructions, please raise your hand if you have any questions. An experimenter will go to your desk and answer your question in private.

During today's session, your payoffs will be in terms of an experimental currency called “experimental pounds”, in short EP. At the end of the experiment, this experimental currency will be converted into British pounds. The amount you earn in this experiment will depend on the decisions that you and other participants make. Your earnings will be paid to you in cash at the end of the experiment. In addition, you will receive £3 for taking part in the experiment.

Please do not talk with others during the session and make sure you have silenced any mobile devices.

**Description of the task**

In this experiment, you will be part of a cohort of 8 participants. The other 7 participants in your cohort can be anyone in this room. Each participant will take on the role of a depositor who has his or her deposits with an experimental ‘bank’. There are 4 banks, named A, B, C and D. You and the other 7 participants will be divided up into 4 groups (2 participants in each group). You will remain in the same group of two and the same cohort for the entire experiment.

The experiment consists of 30 rounds, and in each round your group will be randomly assigned to one of the four banks. At the beginning of each round you and your group-mate will be informed about the bank to which you have been assigned.

At the beginning of each round you and the 7 other persons automatically deposit 100 EP in the bank to which you have been assigned. You must decide whether to withdraw your funds, or to wait and leave your funds deposited with your bank.

In each round one depositor assigned to bank ‘A’ (and bank ‘A’ only) will be randomly chosen and forced to withdraw. Both depositors in bank ‘A’ have an equal chance of being selected and forced to withdraw. If you have been assigned to bank ‘A’, then you will be informed about whether you have been selected and forced to withdraw. If this is the case, the computer will automatically select the action ‘withdraw’ for you. Every other depositor will need to decide whether to withdraw their funds, or to wait and leave them deposited in their bank.

The banks are partially connected to one another as represented in the figure below:

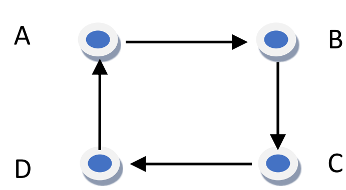

Specifically, banks ‘A’, ‘B’, ‘C’, and ‘D’ are partially connected. Banks are said to be connected to the bank in which they invest part of their deposits. The arrows in the figure display the direction the investment takes place. Here, bank ‘A’ invests in bank ‘B’, which invests in bank ‘C’, which invests in bank ‘D’ which in turn invests in bank ‘A’. So, bank ‘A’ is connected to bank ‘B’; bank ‘B’ is connected to bank ‘C’; and so on. This means that your payoffs depend on your own decision, the decisions of your group-mate, and the decisions of the people in the bank you are connected with. Specifically, how much you earn or lose if you make a withdrawal request or how much you earn or lose by leaving your money deposited in the bank depends on whether your group-mate places a withdrawal request and on how many people in the other bank you are connected with place withdrawal requests. To facilitate your decision, the payoff table below shows the payoffs, that is, the earnings or losses you incur on your 100 EP deposit. The payoff table lists the payoffs that you can obtain depending on your choice, the choice of the other person in your bank, and the choice of the people in the bank you are connected with. Note in the table below that ‘N’ stands for ‘not withdraw’ and ‘W’ stands for withdraw for your choice and the choice of your group-mate. The number of withdrawals in the other, connected bank can be `0W’, `1W’, or `2W’ which stand for 0, 1, or 2 person(s) withdrawing, respectively. Remember that in bank ‘A’ one person is forced to withdraw, so if you are in that bank and you are not forced to withdraw, the column corresponding to no withdrawal request, `N’, by your group-mate is not relevant to you. Also, if you are a depositor in bank ‘D’, the two columns that correspond to zero withdrawal requests, `0W’, in the connected bank are not relevant to you.

| | *Choice of your group-mate* | N | | | W | | |
|---|---|---|---|---|---|---|---|
| | *Number of withdrawals in the connected bank* | 0W | 1W | 2W | 0W | 1W | 2W |
| *Your Choice* | N | 100 | -4 | -83 | -100 | -100 | -100 |
| | W | -15 | -20 | -24 | -32 | -36 | -39 |

Note that since you cannot communicate with others, you must guess what other people will do – whether your group-mate will withdraw (if you are not in bank ‘A’) and how many of the people in bank you are connected with will withdraw (if any) - and act accordingly.

**Procedure**

You will perform the task described above 30 times. Each time is called a round. Each round is completely independent, i.e., you start each round with 100 EP in the bank. At the end of each round, the computer screen will show you your decision and your payoffs for that round. Information for earlier rounds is also provided.

**Computer instructions**

You will see three types of screens: the decision screen, the payoff screen and the waiting screen. Your withdrawal decisions will be made on the decision screen as shown in Figure 1. You can choose to withdraw your funds or leave your funds in the bank by clicking the corresponding option. Note that your decision will be final once you press the 'Confirm' button. The header provides information about what round you are in and the time remaining to make a decision. After the time limit is reached, you will be given a flashing reminder "please reach a decision!"

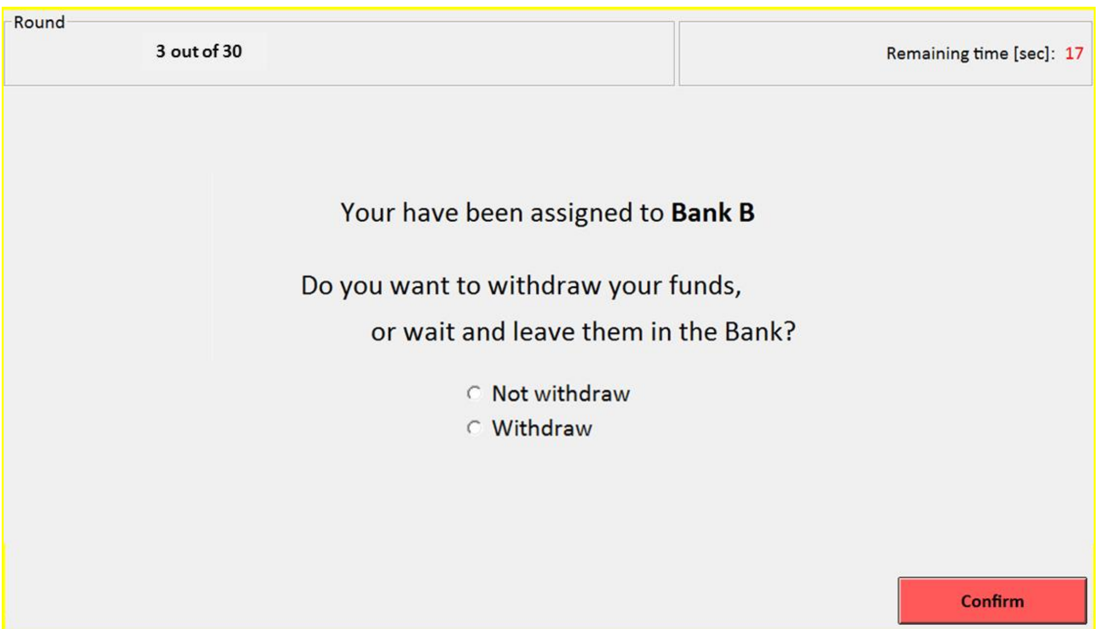


Figure 1: The decision screen

After all participants enter their decisions, a payoff screen will appear as shown in Figure 2. You will see your decision and payoffs for the current round. The history of your decisions, the decisions of your group-mate, the decisions of people in the connected bank and your payoffs is also provided. After you have finished reading this information, click on the "Continue" button to go on to the next round. You will have up to 15 seconds to review the information before a new round begins.

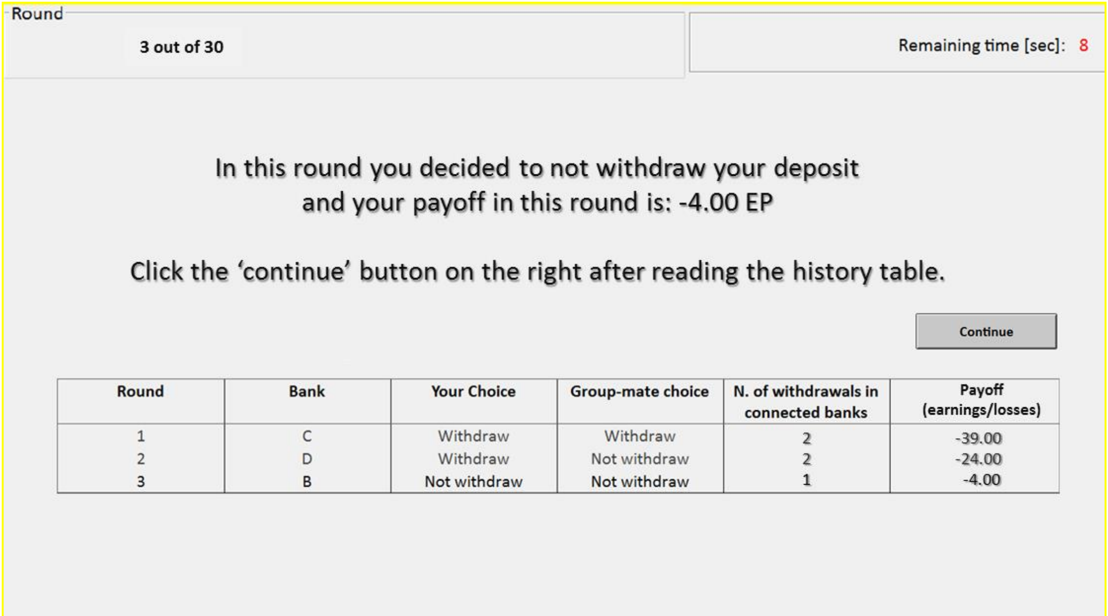


| Round | Bank | Your Choice | Group-mate choice | N. of withdrawals in connected banks | Payoff (earnings/losses) |
|---|---|---|---|---|---|
| 1 | C | Withdraw | Withdraw | 2 | -39.00 |
| 2 | D | Withdraw | Not withdraw | 2 | -24.00 |
| 3 | B | Not withdraw | Not withdraw | 1 | -4.00 |

Figure 2: The payoff screen

You might see a waiting screen (as shown in Figure 3) following the decision or payoff screens. This means that other people are still making decisions or reading information on the outcome of a round and you will need to wait until they finish to go on to the next step.

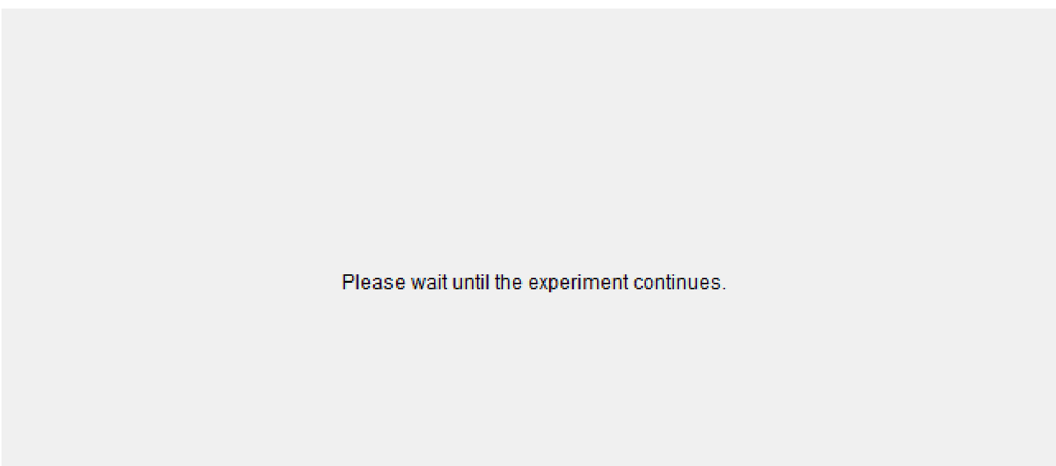


Figure 3: The waiting screen

**Payment**

Once you have completed the 30 rounds the computer program will randomly select 1 round. The payoff (earnings/losses) in the selected round will be added to your deposit of 100EP and transformed into British pounds using the following formula:

[100 EP + Payoff (in EP)] × 0.1

Also, the participation fee will be added to calculate your final earnings. This information will be summarized in your computer screen. After this, you will be asked to answer a short questionnaire. In the meantime the experimenter will prepare your payment. After all participants finish the questionnaire, the experimenter will call you one by one to the payment desk where you will receive your payment in cash.

You may now click start. Before starting we will ask you to complete a comprehension quiz in order to make sure that you understood the instructions. After completing the quiz, you will start round 1.

[Are there any questions?]

## A.2 Printed instructions: Complete network structure

**Instructions**

Welcome to this experiment in economic decision-making. Please pay careful attention to these instructions as they explain how you earn money from the decisions that you make. After we read the instructions, please raise your hand if you have any questions. An experimenter will go to your desk and answer your question in private.

During today's session, your payoffs will be in terms of an experimental currency called "experimental pounds", in short EP. At the end of the experiment, this experimental currency will be converted into British pounds. The amount you earn in this experiment will depend on the decisions that you and other participants make. Your earnings will be paid to you in cash at the end of the experiment. In addition, you will receive £3 for taking part in the experiment.

Please do not talk with others during the session and make sure you have silenced any mobile devices.

**Description of the task**

In this experiment, you will be part of a cohort of 8 participants. The other 7 participants in your cohort can be anyone in this room. Each participant will take on the role of a depositor who has his or her deposit with an experimental 'bank'. There are 4 banks, named A, B, C and D. You and the other 7 participants will be divided up into 4 groups (2 participants in each group). You will remain in the same group of two and the same cohort for the entire experiment.

The experiment consists of 30 rounds, and in each round your group will be randomly assigned to one of the four banks. At the beginning of each round you and your group-mate will be informed about the bank to which you have been assigned.

At the beginning of each round you and the 7 other persons automatically deposit 100 EP in the bank to which you have been assigned. You must decide whether to withdraw your funds, or to wait and leave your funds deposited with your bank.

In each round one depositor assigned to bank 'A' (and bank 'A' only) will be randomly chosen and forced to withdraw. Both depositors in bank 'A' have an equal chance of being selected and forced to withdraw. If you have been assigned to bank 'A', then you will be informed about whether you have been selected and forced to withdraw. If this is the case, the computer will automatically select the action 'withdraw' for you. Every other depositor will need to decide whether to withdraw their funds, or to wait and leave them deposited in their bank.

The banks are fully connected to one another as represented in the figure below:

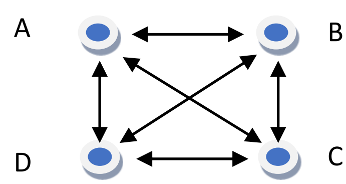


Specifically, banks 'A', 'B', 'C', and 'D' are fully connected. When banks are connected it implies that they invest part of their deposits in the banks they are connected with. The arrows in the figure display the direction the investment takes place. Here, all banks invest in all other banks. This means that your payoffs depend on your own decision, the decisions of the other people in your group, and the decision of the people in the banks you are connected with. Specifically, how much you earn or lose if you make a withdrawal request or how much you earn or lose by leaving your money deposited in the bank depends on whether your group-mate places a withdrawal request and on how many people in the other three banks you are connected with place withdrawal requests. To facilitate your decision, the payoff table below shows the payoffs that is the earnings or losses you incur on your 100 EP deposit. The payoff table lists the payoffs that you can obtain depending on your choice, the choice of the other person in your bank, and the choice of the people in the banks you are connected with. Note, in the table below 'N' stands for 'not withdraw', 'W' stands for withdraw for your choice and the choice of your group-mate. The number of withdrawals in the other, three connected banks can be `0W' '1W', '2W', '3W, '4W', '5W' and '6W' which stand for 0, 1, 2, 3, 4, 5, or all 6 person(s) withdrawing respectively. Remember that in bank 'A' one person is forced to withdraw, so if you are in that bank and you are not forced to withdraw, the column corresponding to no withdrawal request, `N', by your group-mate is not relevant to you. And if you are a depositor in bank 'B', 'C', or 'D', the columns corresponding to zero withdrawal requests in the connected banks are not relevant.

| | *Choice of your group-mate* | N | | | | | | | W | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | *Number of withdrawals in the connected banks* | *0W* | *1W* | *2W* | *3W* | 4W | 5W | 6W | 0W | 1W | 2W | 3W | 4W | 5W | 6W |
| *Your Choice* | N | 100 | 67 | 25 | -18 | -77 | -100 | -100 | -100 | -100 | -100 | -100 | -100 | -100 | -100 |
| | W | -13 | -15 | -17 | -19 | -22 | -24 | -26 | -29 | -31 | -33 | -34 | -37 | -38 | -40 |

Note that since you cannot communicate with others, you must guess what other people will do – whether your group-mate will withdraw (if you are not in bank 'A') and how many of the people in banks you are connected with will withdraw (if any) - and act accordingly.

**Procedure**

You will perform the task described above 30 times. Each time is called a round. Each round is completely independent, i.e., you start each round with 100 EP in the bank. At the end of each round, the computer screen will show you your decision and your payoffs for that round. Information for earlier rounds is also provided.

**Computer instructions**

You will see three types of screens: the decision screen, the payoff screen and the waiting screen. Your withdrawal decisions will be made on the decision screen as shown in Figure 1. You can choose to withdraw your funds or leave your funds in the bank by clicking the corresponding option. Note that your decision will be final once you press the 'Confirm' button. The header provides information about what round you are in and the time remaining to make a decision. After the time limit is reached, you will be given a flashing reminder "please reach a decision!"

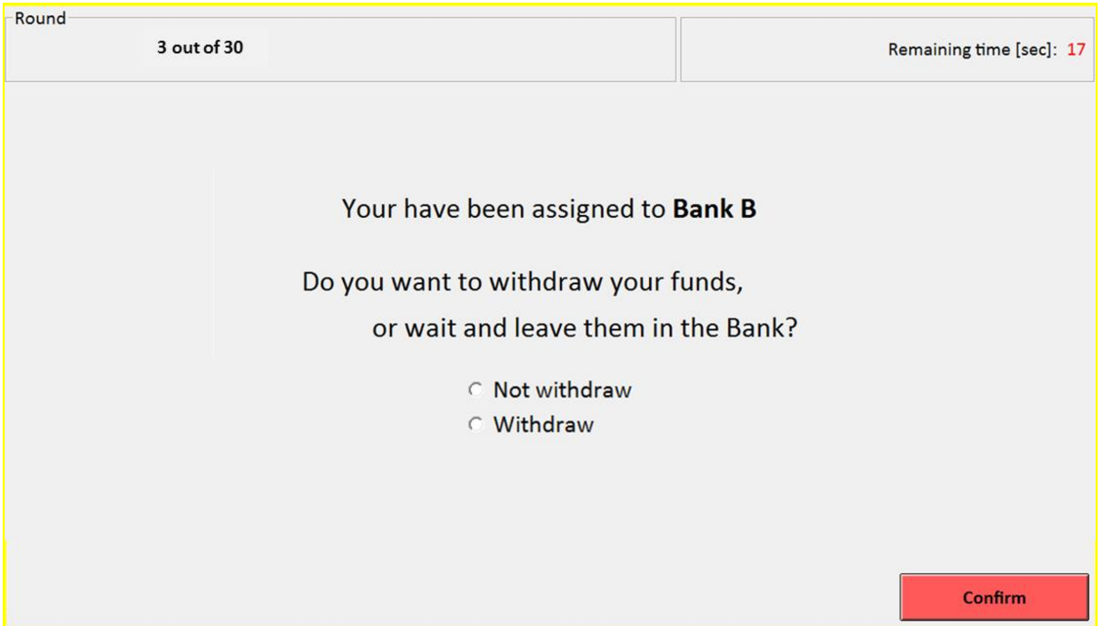


Figure 1: The decision screen

After all participants enter their decisions, a payoff screen will appear as shown in Figure 2. You will see your decision and payoffs for the current round. The history of your decisions, the decisions of your group-mate, the decisions of people in the connected banks and your payoffs is also provided. After you have finished reading this information, click on the "Continue" button to go on to the next round. You will have up to 15 seconds to review the information before a new round begins.

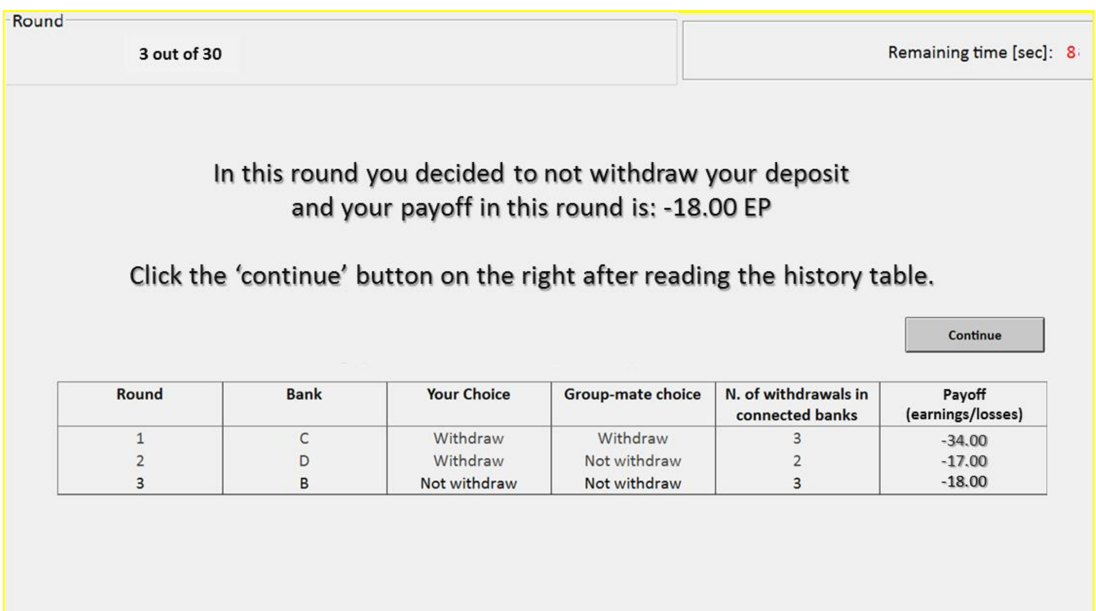


Figure 2: The payoff screen

You might see a waiting screen (as shown in Figure 3) following the decision or payoff screens. This means that other people are still making decisions or reading information on the outcome of a round and you will need to wait until they finish to go on to the next step.

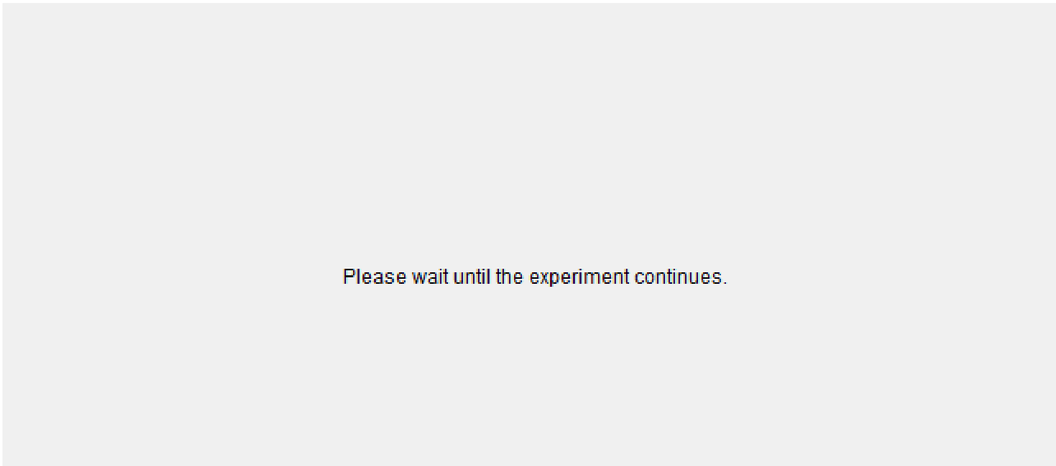


Figure 3: The waiting screen

**Payment**

Once you have completed the 30 rounds the computer program will randomly select 1 round. The payoff (earnings/losses) in the selected round will be added to your deposit of 100 EP and transformed into British pounds using the following formula:

[100 EP + Payoff (in EP)] × 0.1

Also, the participation fee will be added to calculate your final earnings. This information will be summarized in your computer screen. After this, you will be asked to answer a short questionnaire. In the meantime the experimenter will prepare your payment. After all participants finish the questionnaire, the experimenter will call you one by one to the payment desk where you will receive your payment in cash.

You may now click start. Before starting we will ask you to complete a comprehension quiz in order to make sure that you understood the instructions. After completing the quiz, you will start round 1.

[Are there any questions?]