

OTDK Thesis

OTDK-dolgozat

2024

The Effect of Wealth Inequalities on Bank Runs

A vagyoni egyenlőtlenségek hatása a bankcsődök kialakulására



Mihály Fleiner

Corvinus University of Budapest

Economic and Financial
Mathematical Analysis IV

Tamás Vida

Corvinus University of Budapest

Economic and Financial
Mathematical Analysis IV

Supervisor: Miklós Váry, PhD

17th April 2024

Abstract

In this study we investigate the effect of wage inequalities on the risk of bank runs. Empirical evidence suggest that wage and income inequalities can lead to more frequent bank runs.

To better understand this connection we develop an agent-based model (ABM) based on the Diamond-Dybvig bank run model. We run computer simulations with a network of depositors.

Our results suggest that both systematic shocks and wealth inequalities are of great importance to the banking sector, as they increase the risk of a bank run.

Absztrakt

Ebben a dolgozatban a vagyoni egyenlőtlenségek hatását vizsgáljuk a bankcsődök kialakulására. Az empirikus bizonyítékok arra utalnak, hogy a bérek és a jövedelmek egyenlőtlenségei gyakoribb bankcsődökhöz vezethetnek.

Ennek az összefüggésnek a jobb megértése érdekében egy ágensalapú modellt (ABM) dolgozunk ki, amely a Diamond-Dybvig bankcsőd modellen alapul. Számítógépes szimulációkat futtatunk betétesi ismerettségi hálózat használatával.

Eredményeink arra utalnak, hogy mind a szisztematikus sokkok, mind a vagyoni egyenlőtlenségek nagy jelentőséggel bírnak a bankszektor számára, mivel növelik a bankcsőd kockázatát.

Contents

1	Introduction	1
2	Related literature	2
2.1	Diamond-Dybvig model	2
2.2	Behavioural rules and networks	4
2.3	Bank runs and agent-based approaches	5
2.4	Inequality and bank runs	6
3	Model	7
3.1	Investors	8
3.2	Bank	10
3.3	Network	11
4	Simulation	11
4.1	Baseline scenario	12
4.2	Simulation with wealth inequalities	16
4.3	Comparison of the two scenarios	19
4.4	Further analysis of systematic shocks	19
4.5	Further analysis of wealth inequalities	20
4.6	Analysis of the behaviour of the investors	21
4.7	Analysis of the network	23
5	Conclusion	25

1 Introduction

A bankruptcy of a bank is an event in which vast amount of people may loose significant part of their savings in the matter of a few days or weeks. When too many depositors are afraid of loosing their deposits and precisely because of this fact they withdraw their deposits, than they may unintentionally cause a bank run resulting in a bankruptcy. In the fractional-reserve banking system a bank run causes bankruptcy only when sufficiently many depositors withdraw their deposits. In fact it is not the number of depositors who withdraw, but the amount withdrawn that causes the problems. Thus we could assume that greater wealth inequalities could lead to more bank runs, since fewer people could cause one. Our study investigates whether this connection stands.

A cornerstone of bank run literature is the model of Diamond and Dybvig, 1983. The model shows that banks can provide an allocation superior to allocations provided by exchange markets, however even well-functioning banks can face runs, which lead to real economic damage.

We run simulations using an agent-based model where agents represent the heterogeneous investors as well as a bank. The investor agents can be divided into two groups: low-income and high-income agents. We build on the Diamond-Dybvig model, thus our model relies on strong theoretical foundations. With realistic behavioural assumptions we can fine-tune our model to the real world.

Our study adds to the existing literature by providing a realistic, bottom-up model which clearly shows the connection between wealth inequality and bank runs. We also show that systematic shocks - such as a pandemic or a sudden inflation wave - pose serious risk on banks.

A disadvantage of the original Diamond-Dybvig model is being a static model, thus lacks the possibility of inferences on the transition dynamics between a healthy bank and a bank run. Our model extends the original model to a dynamic model, thus we can make inferences on the factors that drive an economy to a bank run.

2 Related literature

2.1 Diamond-Dybvig model

A cornerstone of bank run literature is the model of Diamond and Dybvig, 1983.¹ The model shows that banks can provide an allocation superior to allocations provided by exchange markets, however even well-functioning banks can face runs, which lead to real economic damage.

In this model we assume three periods which we denote 0, 1 and 2. There is a single good in the economy and a possible long term investment in period 0 which in case of a unit investment yields $R > 1$ in period 2 and can be liquidated in period 1 without gaining or loosing on the initial investment. This means that an initial unit investment either results 1 in period 1, or R in period 2.

We assume a continuum of agents, who are homogeneous as of period 0. Each agent has a unit endowment in period 0. The agents can be either type 1 or type 2, which they learn in period 1. (Throughout this paper we might refer to type 1 agents as '*early consumers*' and type 2 agents as '*late consumers*'.) Type 1 agents only value consumption in period 1, while type 2 agents are indifferent between consumption in period 1 and period 2. Thus the utility function conditional on the state of θ can be formulated the following way:

$$U(c_1, c_2) = \begin{cases} u(c_1), & \text{if agent is type 1} \\ \rho u(c_1 + c_2), & \text{if agent is type 2} \end{cases}$$

where $\rho < 1$ is the discount parameter.

In the case of late consumer the authors' interpretation of this type of utility function is that the consumers can '*store*' consumption between periods. Since the long-term investment can be liquidated without a loss in period 1, there is no motivation for either types of agents to not invest their full amount of money in period 0. In autarchy, that is an economy without a bank, neither early, nor late consumers will store between periods, since for early consumers only period 1 consumption matters and late consumers are indifferent between the two consumption, but the technology yields more consumption in period 2. However, as we will see, in the case of a bank, (when agents do not hold the long-term investment directly, but indirectly via a bank) late consumers might have an incentive to store consumption between period 1 and period 2. In fact they will withdraw from the bank in period 1 if they expect a bank run.

¹For our literature review David Romer's Advanced Macroeconomics textbook proved to be very helpful.

Each agent is endowed with one unit of goods which they may deposit in the bank. In our model we loosen this assumption by including poor and rich agents.

Being type 1 or type 2 is publicly unobservable in the model, so it is not possible to write insurance contracts on being type 1. Because of that, in autarchy agents do not interact neither directly nor indirectly. That means each agents invests their endowment in the long term investment. Early agents liquidate their investment in period 1 and late agents hold to their investment until period 2 resulting in R consumption in period 2. We write that by $c_1^1 = 1, c_2^2 = R$ and $c_1^2 = c_2^1 = 0$, where c_i^k denotes the consumption of a type k agent in period i .

If types were publicly observable in period 1, agents could write optimal insurance contracts against the risk of being type 1. From the diminishing marginal utility $c_1^{1*} > 1$, and $c_2^{2*} < R$ follow, with having $c_1^{2*} = c_2^{1*} = 0$. Insuring against the risk of being type 1 would result in higher period 0 expected utility for the agents. It also follows from the assumptions that a late consumer reaches higher ex post utility than an early consumer. That is $u(c_1^{1*}) < pu(u(c_2^{2*}))$. For the detailed derivation of this result we suggest the original article or Romer, David, 2019.

The authors show that via a bank this first best setting is accessible even with publicly unobservable types. Agents can deposit their endowment in the bank in period 0, which is invested in the long term investment by the bank. Each agent can withdraw their money either in period 1 or period 2. In period 1 the bank offers c_1^{1*} for any agent that wishes to withdraw. The remaining funds of the bank are divided equally among the agents who withdraw in period 2. As stated before, late consumers have an incentive to only withdraw in period 2. If only the early consumers choose to withdraw in period 1, than agents withdrawing in period 2 (who are in this case the late consumers) will receive c_2^{2*} . This outcome is the so called 'good equilibrium' in which the economy is in the first best scenario. However, there is also a 'bad equilibrium' or the case of a bank run. If type 2 agents believe that not only type 1 agents withdraw in period 1 then their expectation will be that what they will receive in period 2 will be less than c_2^{2*} . If they expect that their withdrawal in period 2 will be *much* less than c_2^{2*} than they might opt to withdraw in period 1. In this case since every agents withdraws, everyone receives 1. Note that in this case early agents reach the same utility as in autarchy, but late consumers will end up reaching lower utility.

This marks the central result of Diamond and Dybvig. Banks can provide liquidity and in 'good equilibrium' the economy is in the first best case. Yet there is also the possibility of a bank run where average utility is less than in autarchy.

Diamond and Dybvig's work was revolutionary by pointing out these two Nash-equilibria. However, their model does not discuss the microeconomic factors that determine which equi-

librium will be realized. Our study aims to address these questions.

2.2 Behavioural rules and networks

Both empirical and experiment-based research show that there are numerous factors that increase the possibility of a bank run such as heterogeneity of the investors, their social network system, their confidence in banks and their cognitive abilities. In case of fundamental problems, more educated, wealthier and more experienced individuals tend to withdraw money from the bank with a higher possibility. Cognitive abilities also help investors make a better decision. The social network of the investors plays an important role of observability as if an individual investor is able to observe more people withdrawing money, he or she will more likely follow their behaviour. These behavioral characteristics of investors could lead to bank runs even if the bank does not have any real functional problems (Kiss, 2018).

There are numerous studies analyzing the properties of real-life networks and attempting to generate random networks of which properties match the empirical evidence from life. One aspect that should be considered is the small-world feature which occurs in a large number of complex networks. The average path length of a graph is the average distance between two nodes and the clustering coefficient describes the average fraction of pairs of neighbors of a node that are also connected to each other (Wang and Chen, 2003). Previous studies suggest that the average path length in some complex networks are small despite the large number of nodes in the graph and they are clustered, i.e. they have relatively large clustering coefficients (Barabási and Bonabeau, 2003). Watts and Strogatz, 1998 proposes a network generating process that creates a clustered small-world network with small path lengths. If we arrange the vertices of the graph in a circle and create edges between some of the closest nodes, and then rewire some of the links randomly, the formed network will have similar properties that previous researches indicated.

Other studies (Amaral et al., 2000, Barabási and Bonabeau, 2003, Wang and Chen, 2003) draw attention that a powerful indicator of a network is the distribution of its degrees. The degrees of a scale-free network follow a power-law distribution which means that there are some central vertices with a lot of neighbours connected to them but most nodes have way less degrees. These scale-free networks fulfill the small-world property and they occur often in the fields of biology, but acquaintanceship and social networks also follow power-law degree distribution (Barabási and Bonabeau, 2003). For example, Ebel et al., 2002 show that the e-mail networks (where e-mail addresses are the nodes and the e-mails sent are the links between them) are scale-free, highly clustered and fulfill the small-world feature.

The research of Davidsen et al., 2002 considers a basic triangle interaction as a form of evolution creating acquaintance networks as it is based on the idea that often one of our friends introduces us to another friend of him or her. Creating graphs with this assumption reproduces not only the high clustering, but also the small-world and scale-free distribution properties.

The investor agents in our model will be homogeneous considering their education, their general trust of the bank and their cognitive abilities. Hence, we can focus on measuring the effect of wage inequalities on bank failures. However, we arrange the agents on a scale-free network with small-world property considering previous literature on acquaintance graphs so that we can model the effect of the behaviours through local interactions between agents. For readers interested in network science we recommend Barabási, 2016.

2.3 Bank runs and agent-based approaches

The model presented by Diamond and Dybvig, 1983 is widely used to understand the causes of bank runs and to recommend optimal ways of avoiding them. According to the literature review and analysis of Pykälämäki, 2023, one of the main insights of the Diamond-Dybvig model is that using short-term deposits to finance long-term loans can lead to bank failures. Uncertainty and negative information plays a key role in the process of a bank run. It is also possible to approach this model with a game theory perspective: it can be seen as a coordination game where the players are the depositors, and one of the two pure equilibria is a bank run: if a large number of depositors withdraw their money early, the other players' rational choice will be to do the same. This is when a bank failure occurs (Yuan, 2016).

Although the Diamond-Dybvig model clarifies under which conditions the two equilibria happen, it does not discuss the processes and dynamics that move the economy from one equilibrium to the other. A modelling technique to understand the dynamics of the model can be an agent-based approach.

Agent-based models contain agents which are autonomous units that can represent heterogeneity, communication, bounded rationality and learning (Crooks and Heppenstall, 2012). Agents can change their behaviour through adaptive learning techniques within a simulation (Macal and North, 2005) that can help us develop a more realistic model. Contrary to the usual top-down economic models, agent-based models have a bottom-up structure.

Agent-based models were already used previously in order to understand bank runs more. Huang and Huang, 2018 analyzed the probability of bank runs using Artificial Neural Networks (ANNs) which is a method that lets the agents have a cognitive-affective network.

This method allows heterogeneous agents to have slightly different beliefs about the true payoffs. The simulations explain how bank failures can occur even when the bank has no fundamental problems as the agents can react more to the negative announcements about the bank. The authors also recommend a way of avoiding bank runs when the economy starts to seem unstable: banks should increase the late payoffs and so prevent panic and bank failures with the help of the government.

Previous studies also compared different strategies after a bank failure happened. Klimek et al., 2015 examined three mechanisms: to close the bank via a purchase & assumption transaction, to bail-out using the taxpayers' money or to bail in and share the burden among the banks' debtors. Their simulation results showed that under favorable economic conditions, the optimal method is the purchase & assumption transaction, while during distressed economy and recession, it is better to choose the bail-in tool.

Only a few studies so far merged Diamond-Dybvig models with agent-based modelling. Arifovic, 2019 used the individual evolutionary (IEL) algorithm as a tool of adaptive learning. Agents can decide whether they follow a random sunspot announcement about waiting or withdrawing, but they can also develop different strategies using the IEL algorithm. The result suggests that the higher the strategic uncertainty, the more simulations converge to the bad equilibrium. The author also emphasizes that their findings correspond to the ones with human experiments. Santos and Nakane, 2021 created a model of multiple round Diamond-Dybvig economy. Contrary to other studies, the number of banks could change over time and they also applied a sequential service restriction. The model contains a behavioural rule via local interaction: the agents can observe their neighbours and upgrade their strategies with that information gained. To define neighbours, the authors introduce a two-dimensional lattice and Moore neighborhood, i.e. each agent have eight neighbours. The structure of the local interactions have an impact on the frequency of bank runs as well as on the optimal number of banks and the market concentration.

The main goal of our study is to measure the influence of wage inequalities on bank runs for which we apply local interaction among the clients of the bank. For simplicity, we consider only one bank and we run simulations of a multiple round Diamond-Dybvig economy.

2.4 Inequality and bank runs

Previous studies also investigate the effect of inequalities on the financial system and bank runs. Kumhof et al., 2015 argues that two major financial crises of the US -the Great Depression of 1929 and the Great Recession of 2008- were both preceded with periods characterized

with rising income inequalities. They provide a dynamic stochastic general equilibrium model in which greater income inequalities can result in crisis. Amountzias, 2019 provide empirical evidence of 33 OECD countries over 1995-2015 that there is positive relationship between income inequality and financial fragility. Not only income inequality, but wage inequalities also increase the incentives of the run of the bank (Garcia and Panetti, 2022).

Empirical results show that there is correlation between growing inequalities and bank runs. Our paper adds to the existing literature by presenting a model with strong theoretical foundations that shows that inequalities do cause more frequent bank runs. Thus our paper explains the existing empirical evidence by shedding light on the underlying causes of a bank run.

3 Model

Our model applies the original Diamond and Dybvig, 1983 bank run setup that has three periods: 0, 1 and 2. We repeat these three periods after each other, these repetitions we call rounds.

The interpretation of these repeating three-period economy could be that period 0 is the time in the month, when agents receive their wages from their occupation. As in real life wages are hardly ever paid out in cash but almost always transferred to a bank account, our agents do not optimize how much to deposit, they deposit all of their money in the bank. In real life people have their usual everyday expenses, which follow roughly the same trend every month. Banks can plan in advance these 'withdrawals', so this type of behaviour does not pose a bank run threat. In our economy consuming only in period 2, or as we could interpret at the end of the month, is the equivalent of this previously described usual, trend-following consumption. However in real life there are situations where people do not follow their usual consumption trend, because they encounter a special event, such as a repairment of a broken down car, a purchase of an expensive consumer durable, an expensive gift etc. The common feature of these situations is that an extra, during-the-month payment is made. In our model that is the case, when agents consume in period 1.

Since we are applying the Diamond-Dybvig model, our economy also has two equilibria, a 'good' and a 'bad', latter is the case of a bank run. The original model stops here. It notes that there are two, qualitatively different equilibria, but does not say what causes the transition between the two.

Our main focus is the transition between the 'good' and 'bad' equilibrium and the effect of

wealth inequalities on this transition. To answer these questions we put the Diamond-Dybvig model to an ABM setting. We assume a network of investors and behavioural rules both for the investors and for the bank. To include wealth inequalities we also loosen the assumption that every investor has a unit endowment in period 0.

For the results of the original model only a few assumptions of the utility function was required. In our model we need a specific utility function. In Appendix A we show that if we choose CRRA utility functions, than the bank's offer for an investor who invests α in period 0 is α times as much as for an investor who invests a unit. This is a welcome feature as in real life banks usually determine interest rates and pay interest on deposits.

Our model puts t Diamond-Dybvig models after each other where the period 2 is followed by the next period 0 of the next round.

There are two types of agents: one bank and n investors. However, since there is only one bank, often we refer to the investors simply as agent and we refer to the bank as bank. The investors receive income at every period 0. Withdrawing money in period 1 represents early consumption while consuming in period 2 represents late consumption from the original Diamond and Dybvig, 1983 model.

In section 3.1 we describe the variables and parameters as well as the behavioural rules of the investor agents, in section 3.2 we show how the bank establishes the offer in each period and in section 3.3 we present the generating process of the social network of investors.

3.1 Investors

We examine the effect of wage inequality in our research. There are two types of investors: poor and rich ones. Each poor investor receives 1 and each rich one gets α income in every period 0. Both types of investor agents consumes all of their income in one of the next 2 periods, either in period 1 or in period 2. Investors cannot transfer wealth across rounds, so in every period 0 they only have their initial endowment.

We assume that each agent has a CRRA utility function with an impatience parameter ρ :

$$U(c_1, c_2) = \begin{cases} \frac{c_1^{1-\gamma}}{1-\gamma}, & \text{if agent is type 1} \\ \rho \left(\frac{c_1^{1-\gamma}}{1-\gamma} + \frac{c_2^{1-\gamma}}{1-\gamma} \right), & \text{if agent is type 2} \end{cases}$$

The investors have no information in advance whether they will be in need of liquidity at period 2 or they can keep their money in the bank until period 2. They know the expected

probability $E[\theta]$ that they will become early consumers. We assume that each agent has the same expected probability and it is constant over time. In each period each investor agent receives an idiosyncratic shock as a random normally distributed variable with mean 0 and standard deviation σ_{id} . There is a p_{sys} probability of a systematic shock which represents the occurrence of unforeseen liquidity shocks (such as the coronavirus pandemic or the Global Financial Crisis). These shocks are also normally distributed with mean μ_{sys} and standard deviation σ_{sys} . Therefore the actual probability of investor i becoming early consumer at step t is:

$$\theta_{i,t} = E[\theta] + \epsilon_{id,i,t} + \epsilon_{sys,t}$$

where $\epsilon_{id} \sim N(0, \sigma_{id})$ and $\epsilon_{sys} \sim N(\mu_{sys}, \sigma_{sys})$ with probability p_{sys} and 0 with probability $1 - p_{sys}$.

The investors who became early consumers decide to withdraw their money $c_{1,i,t}$ from the bank. Then, every other investor agent observes their neighbours in the social network and recognizes that x proportion of their neighbours withdrew their money. The more neighbours withdrew money, the more becomes the probability of a bank failure and so there is a $p(x)$ probability that the late becomes an early consumer after observing their neighbours, where

$$p(x)_{i,t} = \left(\frac{[x_{i,t} - E[\theta]]_+}{1 - E[\theta]} \right)^\phi.$$

Note how the choice of ϕ affects the probability of becoming an early consumer after interaction. Since the value of the fraction is less than 1, a lower ϕ results in higher probability.

Concluding the description of the investor it is important to note that although we as-

Variable name	Description
$\theta_{i,t}$	Probability of being an early consumer
$Early_{1,i}$	Dummy: 1 if the agent is an early consumer because of $\theta_{i,t}$
$Early_{2,i}$	Dummy: 1 if the agent is an early consumer after interaction

Table 1: Variables of the investor agents

sume utilities for the investors, they are not utility maximizing agents. The investors follow boundedly rational behavioral rules instead of making expectations and maximizing their utility accordingly. The first decision of an investor would be whether to deposit in the bank and if so, how much. In our model we do not have participation constraints or maximization problems, in every period 0 the investor deposits all of their money in the bank. The second

decision would be for late consumers to determine how much to consume in period 1. In our approach, we assume that late consumers are either afraid of a bank run, in this case they withdraw all they can; or they are not afraid and they wait with their consumption until period 2. Being afraid and consuming in period 1 is characterized by the $Early_{2,i}$ dummy variable.

3.2 Bank

There is also a bank in the economy which collects debits in period 0 and offers possible withdrawals or consumptions in period 1 and period 2. As shown by Diamond and Dybvig, banks can offer the social best levels both in period 1 and period 2. In the original model every agent has the same endowment in period 0. In our model we loosen this assumption and we distinguish two types of agents, rich and poor agents. However, we would like to preserve, that the ratio between the possible consumption in period 1 or period 2 and the initial endowment remains equal for all agents. That is as if the bank determined an interest rate in period 0 for money being invested until period 1 and another one for until period 2. We shown in Appendix A, if we assume CRRA utilites for the consumers, this interest rate is independent of the initial period 0 deposit.

In the original model the value of θ , that is the fraction of 'early' consumers is known. This way the bank can calculate and offer the expected utility maximizing offer in period 1 and period 2. Note, that in this setting for a good equilibrium it is required that none of the late consumers consume in period 1. If only a single late consumer would choose to consume in period 1, the bank could not pay the remaining late consumers the previously determined value.

In our model the value of θ is not known. Note that this is not an asymmetric information problem: since late consumers consume in period 1 with a stochastic probability, the fraction of agents who consume in period 1 is the realization of a stochastic process. Thus the bank can only forecast this fraction in period 0. To remain as close to the original model as possible, we assume that the bank uses a profit margin as a puffer. We need to include this imperfection for the bank to be robust against only a few originally late consumer turning to early consumer.

The bank's offer is determined the following way: the bank makes an expectation of the fraction of consumers that will consume in period 1. The bank calculates the first best offers as if θ (fraction of early consumers in the original model) were equal to this expectation and then reduces it with the profit margin. During the calculation of the first best offer, the bank

maximizes the expected utility of the agents. For this the assumed utility of the investors is required.

We assume a heuristic behaviour of the bank. In every round the bank makes an expectation of θ by taking the average of its previous expectation and the realized value of the share of early investors.

For the profit margin puffer we use 3%, but in case the reduced value would be below the initial investment, than the bank's offer is the initial investment.

Concluding the description of the bank we point out that there exists this trade-off between theoretical and applied models. The original model is a purely theoretical model, thus banks offer the first best offers. However, we do not gain much information about what causes a transition between the good and bad equilibrium. Our model on the other hand is a more applied one, thus we have to include imperfections, but doing so we gain more insight on the transition between the two equilibria.

3.3 Network

In our model we assume a network of the agents. This could be interpreted as the acquaintances of the agents. In the real world acquaintances often form scale-free networks with small world property. (Barabási and Bonabeau, 2003) In our baseline scenario we assume such a network. When we introduce wealth inequalities to the model we adjust the network to account for the empirical fact that members of different groups often form clusters in networks. (Ebel et al., 2002) To get some segregation between poor and rich agents we loop through the edges connecting rich and poor agents and with 50% probability we delete the link and redraw it between either two poor or two rich agents. This way poor agents will be over-represented in poor agents neighbours and rich agents will be similarly over-represented in rich agents neighbours.

4 Simulation

We implemented our model in Laboratory for Simulation Development (LSD) software. (Valente, 1998) This is a powerful and easy-to-use platform for various types of agent-based models. For data analysis we used RStudio, for data visualisation we used the ggplot2 R package. (Wickham, 2016)

In each discussed case we assume 1000 investors, 500 rounds and in order to get robust results we run 100 simulations.

Table 2 contains the assumed parameters, unless specified otherwise. Here γ denotes the

Parameter	Value	Parameter	Value
number of investors	1000	number of rounds	500
γ	2.5	R	1.5
exponential scale factor	2.4	average out degree	50
p_{sys}	5%	$(\mu_{sys}, \sigma_{sys})$	(10%, 2%)
σ_{id}	5%	default θ	0.3
ϕ	1		

Table 2: Table of parameters

parameter of the CRRA utility function, R is the return of the long term investment, the exponential scale factor and average out degree are parameters of the network generating algorithm, default θ is the agent's default (before shocked by systematic or idiosyncratic shocks) probability of becoming an early consumer, or changed by the effect of neighbours' actions and ϕ is the parameter in the equation of the probability of an agent's panic. As we can see our model builds on quite a few parameters. This is a usual characteristic of agent-based models. In our analysis we compare different scenarios, thus we can choose the values rather freely, as they mainly remain the same throughout the comparisons.

First in section 4.1 we present a baseline scenario for a benchmark. In section 4.2 we present our model with wealth inequalities.

4.1 Baseline scenario

First we present the baseline scenario, that is an economy with homogeneous investors regarding the initial investment. This is in fact a specific case of our general model, by choosing $\alpha = 1$ both the rich and poor investors have a unit initial investment. In this case the network of the investors is a regular scale-free small-world network.

We run this model as this is the benchmark to the later models which include heterogeneity regarding the initial endowment. When we interpret the results of later models we always compare to these results.

As we simulate 100 simulations with 500 rounds, we have 50.000 rounds in total. Out of

those 1205 bank runs happened, which is around 2.4% of the rounds. This means, that in our model a bank run is a rare event, in more than 97% of the time the bank can serve the investors at the predetermined offer. Figure 1 shows the histogram of the number of bank runs throughout the 100 simulations. As we can see, the number of bank runs is somewhat




Figure 1: Histogram of the number of bank runs per simulation

symmetric with few extremely low or high values. The average number of bank runs per simulation is around 12. We remind that as one simulation consists of 500 rounds this is only 2.4% frequency.

It is interesting to investigate the effect of the systematic shocks. Since systematic shocks are always positive, in case of such a shock, the bank faces higher risk of a bank run. Table 3 contains the contingencies of bank runs and systematic shocks. In the Total row we include the share of bank runs when there is no systematic shock (first column) and when there is systematic shock (second column). In the Total column we calculate the share of systematic shock when no bank run occurred (first row) and similarly the share of systematic shocks when bank runs occurred (second row). The contingency table lets us gain further insight

	No systematic shock	Systematic shock	Total
No bank run	47.309	1486	3.05%
bank run	216	989	82.07%
Total	0.45%	39.96%	

Table 3: Contingency Table of bank runs and systematic shocks

into the nature of bank runs. Without a systematic shock, the event of a bank run is extremely unlikely (0.45%). Although on average there is only 2.4% chance of a bank run, in

case of an systematic shock, the risk of the failure of the bank is almost 40%. As we can see, the large majority of the bank runs occur when a systematic shock is present, however we observed 216 cases, when even without such a shock, the bank failed.

We investigate further whether the systematic shocks indeed cause the bank to fail. Another reason for a bank run could be, when there is panic in the network and agents became early consumers because of their neighbours also panic. In this case we would see a higher fraction of early agents who became early because of their neighbours in case of a bank run.

	Type of early	
	1	2
No bank run	66%	34%
bank run	70%	30%

Table 4: Breakdown of early consumers

Table 4 shows the breakdown of becoming an early consumer. Type 2 of an early agent is an agent, who withdraws early because she panics as her neighbours also withdraw. Being type 1 early agent means being a 'regular' early agent. The table shows that in fact the share of early agents is lower in case of a bank run. This is consistent with our previous result, thus we can infer, that in an economy with perfect wealth equality it is the systematic shocks that cause bank runs not the panic in the network. Later in Section 4.6 we investigate further the topic of panics in the network.

Figure 2 presents the time series of the average offers of the bank. The plot suggests that the




Figure 2: Average values of the bank's offer

offer is staying consistent throughout the rounds. We can see that both c_1 and c_2 are greater

than 1 but lower than R , with $c_2 > c_1$. This is a familiar result from the Diamond-Dybvig model: the bank provides liquidity and risk-sharing among the agents. Figure 3 presents the histograms of the bank's offer side by side. The average value of $c_1 = 1.07$, $c_2 = 1.32$.




Figure 3: Histograms of the bank's offer

We have seen that the number of bank runs is relatively low. This suggests that the bank's expectation of θ and the fraction of early investors (or the 'realization' of θ) progresses very similarly. Figure 4 shows this correlation.




Figure 4: Progression of expected and realized θ

In this section we provided an agent-based extension of the Diamond-Dybvig bank run model. Already from this baseline scenario we could draw important conclusions. We saw that bank runs are a rare event, and that the large majority of bank runs are caused by systematic shocks. In the following sections we advance this model to analyse wealth inequalities and the effect of the network.

4.2 Simulation with wealth inequalities

Now we get to the analysis of the effects of inequalities on bank runs. To investigate this factor, we make two adjustments from the baseline scenario. Firstly, we adjust the network of agents, so that there should be some segregation between the poor and rich agents. Secondly we set $\alpha > 1$ to make the agents heterogeneous in the initial endowment. During this simulation we assume that 80% of the agents have a unit endowment ('poor' agents) and the remaining 20% have an endowment of 5 in every round. ('rich' agents.)

As in the baseline scenario, we run 100 simulations with 500 rounds in each one. Out of these 50.000 rounds we experience 1675 bank runs, which is 3.35% of all cases. Note that in the baseline scenario we had 1205 bank runs which made for only 2.4% of the cases. Introducing wealth inequalities to the model increased the number of bank runs by almost 40%. Figure 5 shows the distribution of bank runs over the simulations. On average 16.75 bank runs




Figure 5: Histogram of the number of bank runs per simulation

occurred in a simulation. Comparing these results to the baseline scenario it is clear that including heterogeneity in terms of initial endowment and adjusting the network to include some segregation lead to more frequent bank runs.

Looking at the contingency table of bank runs and systematic shocks confirms our preliminary expectations. Table 5 contains the breakdown of bank runs and systematic shocks, with the values of the baseline scenario in brackets.

	No systematic shock	Systematic shock
No bank run	46.930 (47.309)	1395 (1486)
Bank run	595 (216)	1080 (989)

Table 5: Contingency Table of bank runs and systematic shocks

The greatest change is to be seen in the case of a bank run without a systematic shock. We remind that in the baseline scenario 82% of the bank runs occurred when a systematic shock was present. In the current setting the share of bank runs without a systematic shock increased notably, from being less than one fifth of the cases to becoming more than one third of the occurrences.

As we noted in the baseline scenario, the share of panic driven early agents were not responsible for the bank run. In the current setting we get similar results regarding the breakdown of early agents. Table 6 shows the fraction of early agents who became early due to panic (type 2 early agent). In brackets the results of the baseline scenario is included for easier

	Type of early	
	1	2
No bank run	66% (66%)	34% (34%)
bank run	69% (70%)	31% (30%)

Table 6: Breakdown of early consumers

comparison. As we can see, the table is almost identical with the table in the baseline scenario. This result suggests that rich investors do not cause bank runs by starting panic waves among each other, but by withdrawing independently of each other.

To gain further insight into the dynamics of a bank run, we investigate share of rich agents among the early agents in case of with and without a bank run. This breakdown is shown in Table 7. From the table we can tell that in case of a bank run 64% of the rich agents are

	Share of early	
	poor	rich
No bank run	56%	56%
bank run	61%	64%

Table 7: Share of early agents among poor and rich agents

early consumers. When the bank is healthy, there is no difference between the fraction who becomes early consumer among the rich and poor agents. However, as we see in the second row of the table, during bank runs rich agents are over-represented among the early agents. This strengthens our initial hypothesis that higher wealth inequalities lead to higher risk of a bank run.

Figure 6 shows the average values of the bank's c_1 and c_2 offers. Once again, the offers




Figure 6: Average values of the bank's offer

remain consistent over rounds. Also the unfolding of the offer is very similar as we have seen in the baseline scenario.

Figure 7 shows the histograms of the bank's c_1 and c_2 offer. Both the distribution and the




Figure 7: Histograms of the bank's c_1 and c_2 offers

average values remain similar as in the baseline scenario. The average value of $c_1 = 1.07$, $c_2 = 1.32$.

Lastly, Figure 8 presents the connection between the bank's expectation of the share of early agents and the realized value of early agents.




Figure 8: Progression of expected and realized θ

4.3 Comparison of the two scenarios

We ran two scenarios to investigate the effect of wealth inequalities on bank runs. In the baseline scenario every agent was homogeneous in terms of the initial investment and the agents formed a scale-free small-world network. In the second setting 20% of the agents had 5 units to invest while the remaining 80% had a unit endowment. The network structure was slightly modified to represent the fact, that members of the same group are more likely to know each other than members of different groups.

In the baseline scenario the probability of a bank run was 2.4%. That grew by almost 40% to 3.35% in the model containing inequalities. The number of bank runs when there were no systematic shock more than doubled in the latter setting. Table 7 also suggests that the presence of rich agents is the factor that causes bank runs.

Our model proves our initial hypothesis that higher wealth inequalities lead to higher risk of a bank run. Our results also suggest that the channel which this effect plays through is not that there would be panic shocks in the circles of rich agents, but that systematic and/or idiosyncratic shocks occur so, that rich agents become early investors which leads to the failure of the bank.

4.4 Further analysis of systematic shocks

As we have seen systematic shocks are the main reason of bank runs in the baseline scenario. To further investigate the robustness of this effect we run separate simulations, where we vary the probability of a systematic shock p_{sys} and the size of the shock μ_{sys} .

To get neat results we run the simulation in the simpler baseline scenario. We run simulations

with 4 different setting. The 4 setting which we simulate are the followings: $(p_{sys}, \mu_{sys}) \in \{(5\%, 5\%), (5\%, 10\%), (10\%, 5\%), (10\%, 10\%)\}$

Table 8 contains the number of bank runs in each setting. As anticipated, the lowest number

	$\mu_{sys} = 5\%$	$\mu_{sys} = 10\%$
$p_{sys} = 5\%$	492	1205
$p_{sys} = 10\%$	654	2071

Table 8: Number of bank runs in each setting

of bank runs occurred when both the probability of a systematic shock and the size of a shock was small and the highest number occurred when both of these parameters were higher. Also both parameters on their own have positive correlation with the risk of a bank run.

As previously seen both in the baseline scenario and in the wealth inequality extended model the majority of bank runs happened when a systematic shock occurred. In conclusion we can infer that systematic shocks increase the probability of a bank run and that more frequent or larger shocks also cause more runs. Risk managers in banks and the supervisors working in central banks should be familiar with this effect to achieve a banking sector resistant to systematic shocks.

4.5 Further analysis of wealth inequalities

Our main result is the positive correlation between wealth inequalities and the risk of a bank run. In this section we aim to further strengthen this argument by running different scenarios with different types of wealth inequalities. As in the previous section, we vary two factors, the share of rich agents and the amplitude of how rich the rich agents are, which we do via the α parameter. We run models with the share of rich agents being 10% or 20% and with $\alpha = 5$ and $\alpha = 10$.

Table 9 contains the number of bank runs in each setting. The data suggest that if the rich agents are ceteris paribus richer, - that is the the wealth inequalities are greater - the number of bank runs also increase. An additional result is that the share of rich agents does not seem to have a clear effect when raised from 10% to 20%.

	Share of rich	
	10%	20%
$\alpha = 5$	1665	1633
$\alpha = 10$	2187	2286

Table 9: Number of bank runs in each setting

4.6 Analysis of the behaviour of the investors

In Section 4.2 we have seen that the reason for the growing number of bank runs were not due to panic in the clusters of rich agents. An explanation for that might be that the assumed ϕ parameter in the behavioural rule of the agents were too large, thus the probability of becoming an early investor in the second step were low. In this section we investigate the adjustment of this behavioural rule via adjusting the value of ϕ . We run simulations where $\phi \in \{0.25, 0.5, 0.75, 1\}$.

Table 10 presents the number of bank runs with different values of ϕ . A clear connection between ϕ and the number of bank runs is not present, too low or too high values of ϕ might result in lower number of bank runs though. When the chosen value is too large, the probability of withdrawing money from the bank due to the fear of a bank run is low, thus we rule out these types of bank runs. When ϕ is too small, the economy converges to an equilibrium, where the bank's expectation of θ and the share of early agents are too high. The different values of ϕ might cause different emerging dynamics, thus we investigate these

ϕ	Number of bank runs
0.25	1546
0.5	1887
0.75	1892
1	1627

Table 10: The number of bank runs in each simulation

simulations separately.

In the case of $\phi = 0.25$ the expected θ stabilizes in a relatively high level. Figure 9 shows the distribution of this expectation. The average value is $\bar{\theta} = 75\%$. From the high value of expected θ a low value for c_1 (that is the bank's offer in period 1) follows. Figure 10 presents the distributions of the bank's offers. While c_2 remains relatively similar to previous





Figure 9: Distribution of the bank's expectation of θ

scenarios, c_1 is much lower due to the high number of early agents. Because of the low number



(a) Histogram of c_1 offer

(b) Histogram of c_2 offer

Figure 10: Histograms of the bank's c_1 and c_2 offers

of late agents this setting is not too relevant for our research. This showcases the fact that right parametrization is crucial in these type of models.

The case of $\phi = 0.5$ results in relatively similar results as the just described scenario. In this case average expected θ is $\bar{\theta} = 66\%$ and the average of the bank's period 1 offer is $\bar{c}_1 = 1.05$. The simulations with $\phi = 0.75$ and $\phi = 1$ has sufficiently many late agents to forward with the analysis. Table 11 consists the share of agents who are 'type 2' early. These are the agents who were late originally, but after the interaction with their neighbours they became early agents. We can state that the proportion of agents who become early due to the panic of a bank run is independent of bank runs. This means that we do not avoid this counterintuitive result by simply adjusting the value of ϕ . The reason why we still get this result after adjusting the value of ϕ lies in the structure of the model. When looking at the data, we see that in the case of $\phi = 0.25$, the first few rounds of the simulations are always

	No bank run	Bank run
$\phi = 0.75$	23%	24%
$\phi = 1$	19%	19%

Table 11: Share of type 2 early agents

characterized by bank runs and a large share of agents who are type 2 early. However after a few rounds, the bank's expectation of θ stabilizes in a relatively high value. In contrast agent's still have the same default probability of becoming an early agent and the distribution of idiosyncratic and systematic shocks also remain the same. This together means that the bank can plan with agents panicking, thus it avoids runs.

4.7 Analysis of the network

We would like to investigate further the reasons behind the fact that so far our model does not replicate successfully the empirical experience that panicking agents can cause bank runs. An important part of how panic can spread is the network of the agents. The factor that drives whether an agent becomes early only after interacting with their neighbours is the share of early neighbours. Thus if in period 1 an agent realizes that she was an early agent, this factor also influences her neighbour's probability of becoming an early agent. This effect of the individual's is stronger if an agent has less neighbours. Thus it might be insightful to run simulations with networks with different density.

In our network generating algorithm the average out degree parameter sets the average number of neighbours. We run simulations with this parameter being 10, 25, 50 and 75. Table 12 contains the number of bank runs in each simulation. The effect of the average

Average out degree	Number of bank runs
10	3563
25	2058
50	1760
75	1763

Table 12: The number of bank runs in each simulation

number of neighbours seems clear (or at least in the lower regions), namely having fewer

neighbours lead to more frequent bank runs.

Table 13 presents the share of type 2 early agents in cases of with and without a bank run. When the number of neighbours in the network is the lowest, 27% of all the agents are

Average out degree	No bank run	Bank run
10	24%	27%
25	24%	24%
50	23%	24%
75	23%	24%

Table 13: Share of type 2 early agents

type 2 early, that is they withdraw money only because of they fear a bank run. In normal conditions this share is 3 percentage point lower. This might not seem significant, but note that these are the results of 50.000 simulations, thus we can draw conclusions from smaller differences (3%) too. However this effect vanishes almost entirely as we increase the number of connections in the network. This is no surprise, in fact it follows from the Central Limit Theorem. We remind the equation of becoming a type 2 early agent:

$$p(x)_{i,t} = \left(\frac{[x_{i,t} - \theta_{i,t}]_+}{1 - \theta_{i,t}} \right)^\phi$$

where x is the proportion of the agent's neighbours who became type 1 early and $\theta_{i,t}$ is the agent's own probability of becoming a type 1 early agent. If the agent has more neighbours then the share of early agents among her neighbours is more likely to stay close to its average, thus there will not be emerging panic waves which would cause bank runs.

Because of these considerations we continue our analysis with the case when average out degree is 10. We are interested in whether we see that rich agents' panic cause bank runs. To test these hypothesis we compare the share of rich agents who become type 2 early consumers during bank runs with the similar share of poor agents. Table 14 contains the comparison. Again we see slight but positive difference in favor of the rich agents.

The share of type 2 early agents during bank runs	
poor	rich
26%	28%

Table 14: The share of poor and rich type 2 early agents during bank runs

This concludes our analysis. To summarize, we saw small but persistent results suggesting that panicking rich agents might bank runs. However, this result is not as clear and convincing as our previous results, thus this question should be a topic of further research.

5 Conclusion

In this paper we provided an agent-based extension of the Diamond-Dybvig bank run model. While the Diamond-Dybvig model is a top-down equilibrium model, our model is a bottom-up structured agent-based model, thus it does not require strong equilibrium assumption. This allows us to model emerging non-equilibrium dynamics. Our model sheds light on the transition dynamics between the proper functioning of a bank and the emergence of a bank run.

In our research we mainly focused on the effect of wealth inequalities on bank runs. We provided sound evidence that higher wealth inequalities lead to higher risk of bank runs. This explains the existing literature on the empirical evidence of the connection between bank runs and inequalities.

We presented the clausal connection between systematic shocks and bank runs. Since the banking sector has a key role in the economy, understanding better the risks surrounding it is of great importance. Our results suggest that systematic shocks and wealth inequalities cannot be neglected.

We investigated whether those are the rich depositors who cause bank runs by panicking and spreading panic waves among each other. Our results suggest that this might be an existing factor, but the model did not provide clear and unquestionable proof. Further research into this question should be topic of further research.

References

- Amaral, L. A. N., Scala, A., Barthelemy, M., & Stanley, H. E. (2000). Classes of small-world networks. *Proceedings of the national academy of sciences*, 97(21), 11149–11152.
- Amountzias, C. (2019). An investigation of the effects of income inequality on financial fragility: Evidence from organization for economic co-operation and development countries. *International Journal of Finance & Economics*, 24(1), 241–259. <https://doi.org/https://doi.org/10.1002/ijfe.1660>
- Arifovic, J. (2019). Evolution of sunspot like behavior in the agent based economies of bank runs. *Journal of Evolutionary Economics*, 29, 365–389.
- Barabási, A.-L. (2016). *Network science*. Cambridge University Press.
- Barabási, A.-L., & Bonabeau, E. (2003). Scale-free networks. *Scientific american*, 288(5), 60–69.
- Crooks, A. T., & Heppenstall, A. J. (2012). Introduction to agent-based modelling. In A. J. Heppenstall, A. T. Crooks, L. M. See, & M. Batty (Eds.), *Agent-based models of geographical systems* (pp. 85–105). Springer Netherlands. https://doi.org/10.1007/978-90-481-8927-4_5
- Davidson, J., Ebel, H., & Bornholdt, S. (2002). Emergence of a small world from local interactions: Modeling acquaintance networks. *Physical review letters*, 88(12), 128701.
- Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3), 401–419. Retrieved February 28, 2024, from <http://www.jstor.org/stable/1837095>
- Ebel, H., Mielsch, L.-I., & Bornholdt, S. (2002). Scale-free topology of e-mail networks. *Physical review E*, 66(3), 035103.
- Garcia, F., & Panetti, E. (2022). Wealth inequality, systemic financial fragility and government intervention. *Econ. Theory*.
- Huang, W., & Huang, Q. (2018). Connectionist agent-based learning in bank-run decision making. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(5).
- Kiss, H. J. (2018). Depositors’ behaviour in times of mass deposit withdrawals. *Financial and Economic Review*, 17(4), 95–111.
- Klimek, P., Poledna, S., Farmer, J. D., & Thurner, S. (2015). To bail-out or to bail-in? answers from an agent-based model. *Journal of Economic Dynamics and Control*, 50, 144–154.
- Kumhof, M., Rancière, R., & Winant, P. (2015). Inequality, leverage, and crises. *American Economic Review*, 105(3), 1217–45. <https://doi.org/10.1257/aer.20110683>

- Macal, C., & North, M. (2005). Tutorial on agent-based modeling and simulation. *Proceedings of the Winter Simulation Conference, 2005.*, 14 pp.-. <https://doi.org/10.1109/WSC.2005.1574234>
- Pykälämäki, T. (2023). Understanding bank runs with the diamond-dybvig model.
- Romer, David. (2019). *Advanced macroeconomics* (5th). McGraw-Hill Education.
- Santos, T. R. E. d., & Nakane, M. I. (2021). Dynamic bank runs: An agent-based approach. *Journal of Economic Interaction and Coordination*, 16(3), 675–703.
- Valente, M. (1998). *Laboratory for simulation development*. <https://ideas.repec.org/p/aal/abbswp/98-5.html>
- Wang, X. F., & Chen, G. (2003). Complex networks: Small-world, scale-free and beyond. *IEEE circuits and systems magazine*, 3(1), 6–20.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *nature*, 393(6684), 440–442.
- Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. <https://ggplot2.tidyverse.org>
- Yuan, B. (2016). The diamond–dybvig model of bank runs as a coordination game.

Appendix A

Assuming CRRA utility functions the expected utility of an investor who invests a unit in period 0 is the following:

$$E[u(c_1, c_2)] = \theta \frac{c_1^{*1-\gamma}}{1-\gamma} + \rho(1-\theta) \frac{c_2^{*1-\gamma}}{1-\gamma}$$

where following the usual notation γ denotes the parameter of the utility function, R is the return of the long term investment in the Diamond-Dybvig model and θ is the fraction of early investors. Since in good equilibrium the bank offers the first best offer, we know that only the early consumers withdraw in the first period and late consumers will receive the resulting funds. This way we can write c_2^* as a function of c_1^* :

$$c_2^* = \frac{(1 - \theta c_1^*)R}{1 - \theta}$$

The expected utility takes the following form:

$$E[u(c_1, c_2)] = E[u(c_1)] = \theta \frac{c_1^{*1-\gamma}}{1-\gamma} + \rho(1-\theta) \frac{\left(\frac{(1-\theta c_1^*)R}{1-\theta}\right)^{1-\gamma}}{1-\gamma}$$

From this the utility maximizing c_1^* :

$$c_1^* = \frac{\frac{R^{1-\frac{1}{\gamma}}}{1-\theta}}{1 + \frac{\theta R^{1-\frac{1}{\gamma}}}{1-\theta}}$$

For a rich investor who has α endowment in period 0, the expected utility is the following:

$$E_{rich}[u(C_1, C_2)] = \theta \frac{C_1^{*1-\gamma}}{1-\gamma} + \rho(1-\theta) \frac{C_2^{*1-\gamma}}{1-\gamma}$$

We have that

$$C_2^* = \frac{(\alpha - \theta C_1^*)R}{1 - \theta}$$

We can rewrite the expected utility function once again:

$$E_{rich}[u(C_1, C_2)] = E_{rich}[u(C_1^*)] = \theta \frac{C_1^{*1-\gamma}}{1-\gamma} + \rho(1-\theta) \frac{\left(\frac{(\alpha - \theta C_1^*)R}{1-\theta}\right)^{1-\gamma}}{1-\gamma}$$

C_1^* maximizing the expected utility is the following:

$$C_1^* = \alpha \cdot \frac{\frac{R^{1-\frac{1}{\gamma}}}{1-\theta}}{1 + \frac{\theta R^{1-\frac{1}{\gamma}}}{1-\theta}} = \alpha \cdot c_1^*$$

From here $C_2^* = \alpha c_2^*$ follows.