
MŰHELYTANULMÁNYOK

DISCUSSION PAPERS

MT-DP – 2014/29

**Correlated observations, the law of small
numbers and bank runs**

GERGELY HORVÁTH – HUBERT JÁNOS KISS

Discussion papers
MT-DP – 2014/29

Institute of Economics, Centre for Economic and Regional Studies,
Hungarian Academy of Sciences

KTI/IE Discussion Papers are circulated to promote discussion and provoke comments.
Any references to discussion papers should clearly state that the paper is preliminary.
Materials published in this series may subject to further publication.

Correlated observations, the law of small numbers and bank runs

Authors:

Gergely Horváth
Department of Economic Theory
Universität Erlangen-Nürnberg
email: gergely.horvath@fau.de

Hubert János Kiss
research fellow
'Momentum' Game Theory Research Group
Institute of Economics,
Centre for Economic and Regional Studies
Hungarian Academy of Sciences
Eötvös Loránd University - Department of Economics
email: kiss.hubert.janos@krtk.mta.hu

November 2014

ISBN 978-615-5447-48-8
ISSN 1785 377X

Correlated observations, the law of small numbers and bank runs

Gergely Horváth – Hubert János Kiss

Abstract

Empirical descriptions and studies suggest that generally depositors observe a sample of previous decisions before deciding if to keep their funds deposited or to withdraw them. These observed decisions may exhibit different degrees of correlation across depositors. In our model depositors are assumed to follow the law of small numbers in the sense that they believe that a bank run is underway if the number of observed withdrawals in their sample is high. Theoretically, with highly correlated samples and infinite depositors runs occur with certainty, while with random samples it needs not be the case, as for many parameter settings the likelihood of bank runs is less than one. To investigate the intermediate cases, we use simulations and find that decreasing the correlation reduces the likelihood of bank runs, often in a non-linear way. We also study the effect of the sample size and show that increasing it makes bank runs less likely. Our results have relevant policy implications.

Keywords: bank runs, law of small numbers, samples, threshold decision rule

JEL classification: D03, G01, G02

Korreláló megfigyelések, a kis számok törvénye és a bankrohamok

Horváth Gergely – Kiss Hubert János

Összefoglaló

Empirikus megfigyelések és tanulmányok azt mutatják, hogy általában a betétesek megfigyelik a korábbi döntések egy részét, mielőtt arról döntenének, hogy kivegyék-e a pénzüket a bankból. Ezek a megfigyelések különböző mértékben korrelálhatnak a betétesek között. A modellben azt tesszük fel, hogy a betétesek a kis számok törvényét követik abban az értelemben, hogy azt hiszik: bankroham van kialakulóban, ha a megfigyeléseik között sok a betét kivétel. Elméletileg, végtelen betétes mellett a nagymértékben korreláló megfigyelések biztosan bankrohamhoz vezetnek, míg véletlen megfigyelések mellett a bankroham valószínűsége kisebb, mint egy. A köztes eseteket szimuláció segítségével elemezzük, és azt találjuk, hogy a csökkenő korreláció mérsékli a bankrohamok kialakulásának valószínűségét, gyakran nem-lineáris módon. A megfigyelt minta nagyságának a hatását is vizsgáljuk, és azt találjuk, hogy nagyobb minták csökkentik a bankroham valószínűségét.

Tárgyszavak: bankroham, kis számok törvénye, minta, küszöb döntési szabály

JEL kód: D03, G01, G02

Correlated observations, the law of small numbers and bank runs

Gergely Horváth

Friedrich-Alexander-Universität Erlangen-Nürnberg*

Hubert János Kiss

Eötvös Loránd University and MTA KRTK[†]

October 31, 2014

Abstract

Empirical descriptions and studies suggest that generally depositors observe a sample of previous decisions before deciding if to keep their funds deposited or to withdraw them. These observed decisions may exhibit different degrees of correlation across depositors. In our model depositors are assumed to follow the law of small numbers in the sense that they believe that a bank run is underway if the number of observed withdrawals in their sample is high. Theoretically, with highly correlated samples and infinite depositors runs occur with certainty, while with random samples it needs not be the case, as for many parameter settings the likelihood of bank runs is less than one. To investigate the intermediate cases, we use simulations and find that decreasing the correlation reduces the likelihood of bank runs, often in a non-linear way. We also study the effect of the sample size and show that increasing it makes bank runs less likely. Our results have relevant policy implications.

JEL codes: D03, G01, G02

Keywords: bank runs, law of small numbers, samples, threshold decision rule

*Corresponding author. Department of Economic Theory, Friedrich-Alexander-Universität Erlangen-Nürnberg, Lange Gasse 20, D-90403 Nuremberg, Germany. E-mail: horvathgergely@gmail.com. Phone: +49 (0)911 5302 224.

[†]Department of Economics, Eötvös Loránd University, 1117 Pázmány sétány 1/a, Budapest, Hungary. E-mail: hubert.kiss@tatk.elte.hu. Phone: +36-30-4938062. Kiss is a research fellow in the “Momentum” (LD-004/2010) Game Theory Research Group in the Institute of Economics Centre for Economic and Regional Studies of the Hungarian Academy of Sciences (MTA KRTK).

1 Introduction

Although bank runs were very rare phenomena in developed countries in the decades before 2007, the run on Northern Rock, an English bank, heralded that "... [T]he age of the bank run has returned." (Tyler Cowen, The New York Times March 24, 2012) Episodes of banks and other financial institutions suffering sudden and massive withdrawals of deposits and other funding sources were frequent during the recent financial crisis (e.g. the investment bank Bear Stearns in the US, the DSB Bank in the Netherlands or Bankia in Spain). Deteriorating fundamentals are a prime cause of bank runs, but there was often also a substantial self-fulfilling component to the behavior of depositors. Depositors hurry to withdraw fearing that other depositors' withdrawal will cause the bank to fail. This idea is illustrated vividly by the words of Anne Burke, a client of Northern Rock, who said the following while queuing up to withdraw her funds: "It's not that I disbelieve Northern Rock, but everyone is worried and I don't want to be the last one in the queue. If everyone else does it, it becomes the right thing to do."¹ The above quote shows that depositors react to other depositors' observed decisions. Empirical studies (Kelly and O Grada, 2000; Starr and Yilmaz, 2007; and Iyer and Puri, 2012) and experimental findings (Garratt and Keister, 2009; Kiss et al. 2012) support this idea as well.

It is natural to ask: what do depositors observe? In some cases observability of other depositors' actions is almost non-existent, as it was the case during the silent run on Washington Mutual in 2008, when depositors withdrew their funds electronically.² Other empirical observations suggest two things. First, although it is easier to observe somebody queuing up to withdraw, depositors may also know that others have decided to keep their money deposited. Kelly and O Grada (2000) and Iyer and Puri (2012) point at the importance of observing decisions of both sorts (withdrawal or keeping the money in the bank) in one's social network. Starr and Yilmaz (2007) argue that during a bank run incident in 2001 in Turkey small and medium-sized depositors seemed to observe only withdrawals of their peers but the behavior of large depositors appeared to be driven by observing both choices. Second, the previous empirical studies and other descriptions suggest that not all previous decisions can be observed, depositors observe only a sample of earlier choices.

A noteworthy aspect of the above empirical studies is that none of the banks affected by the runs were fundamentally bad banks. Thus, the massive withdrawals cannot be explained by the decisions of informed depositors withdrawing from a financial intermediary due to fundamental reasons, but a coordination failure

¹See <http://www.bloomberg.com/apps/news?pid=newsarchive&sid=aeyCkzcRIU4>

²When nothing is observed, Arifovic et al. (2013) show experimentally that bank runs are more likely the more stringent are the conditions for the coordination of depositors. Relatedly, Sultanum (2014) characterizes the direct mechanism that implements the constrained optimal outcome in a Diamond-Dybvig setup with aggregate uncertainty and simultaneous decisions.

among the depositors seems to be behind these runs. It is of first-order importance to understand what may cause these coordination failures since it is clearly not optimal that healthy banks are run by depositors and the financial intermediation is disrupted. Although fundamentally weaker banks are more likely to be run by depositors, there is clear evidence (e.g. Davison and Ramirez, 2014; De Graeve and Karas, 2014) that even fundamentally healthy financial intermediaries suffer bank runs. In this study we attempt to further our understanding by assuming a healthy bank that has depositors deciding in a sequential manner and who act upon observing a sample of previous decisions. The main question that we study is how sampling affects the emergence of bank runs. In close-knit communities (for instance, among customers of a rural bank in a village) the information observed by subsequent depositors generally overlaps, depositors who decide sequentially observe to a large degree the same previous choices. On the contrary, for clients of a big bank information is less correlated, depositors deciding consecutively are less likely to observe the same previous decisions. Do these differences lead to a different likelihood of bank run?

To answer the question we develop a model in which depositors observe a sample of the previous actions and decide sequentially. For simplicity, we consider the case when depositors observe samples of the same size. We change the degree of correlation across subsequent samples. A key issue is how depositors extract information from the observed sample. It is not obvious what conclusions a depositor can draw when observing a set of previous decisions that need not be representative and knowing that her decision will be possibly observed by others and hence influence if a bank run starts or not. We use a well-known regularity of decision-making documented in behavioral economics: the law of small numbers.³ It is a faulty generalization, the observed sample is taken to be representative of the population of interest. Hence, if a depositor observes many withdrawals in her sample, then she believes that the number of withdrawals in general will be high so withdrawing early may seem the optimal decision since otherwise the depositor may receive nothing from the bank. On the contrary, when only a few withdrawals are observed, it may suggest that there is no bank run underway, so there is no reason to withdraw the money from the bank. The law of small numbers relies on the availability and representativeness heuristics and leads to overinference, that is the belief that even small samples reflect the properties of the population.⁴

We study analytically two focal cases. In the random sampling case, depositors observe any of the previous decisions with equal probability, so correlation across consecutive samples is low. With overlapping samples, depositors observe the last decisions. This sort of sampling captures several intuitive ideas. On the one hand, it can be argued that recent actions are more likely to be observed. On the other hand, depositors

³For more details, see Rabin (2002) and DellaVigna (2009). Barberis and Thaler (2003) provide related examples in finance.

⁴Tversky and Kahneman (1973) and Kahneman and Tversky (1974) are two important studies on the availability and the representativeness heuristics.

with the same characteristics (those living in the same neighborhood as in Kelly and O Grada (2000) or the depositors grouped according to their deposit size as in Starr and Yilmaz (2007)) tend to observe the same information. Our overlapping sampling is a way to have depositors observe highly correlated information. The intuition why different sampling mechanisms lead to different results is the following. With overlapping samples subsequent depositors observe very similar samples, so if a depositor observes a sample that makes her withdraw, the next depositor will observe at least as many withdrawals as the last one. If depositors follow identical decision rules, then it leads to a bank run. With random sampling the correlation across samples is considerably lower, so subsequent decisions will be less uniform resulting in less bank runs.

Investigating the intermediate cases analytically is too complicated, so we analyze them through simulations, varying the parameters of the model. We find systematically that the correlation of the samples affects the likelihood of bank runs, higher correlation leading to an increased likelihood of bank runs. As we decrease the correlation of samples, the likelihood of bank runs decreases in a non-linear way, suggesting a sharp transition from a state where bank runs are very likely to another where they do not occur. In the simulations we find also that as the sample size increases, the probability of a bank run decreases.

Our results suggest some ways that policy-makers could prevent bank runs that are unjustified fundamentally and are triggered by overly panicky depositors. By providing more information about previous decisions (that is, increasing the sample size) and by attempting to make that information the least correlated, the likelihood of bank runs can be considerably diminished. The threshold in the decisions rule of the depositors can also be affected by credible policies and safety network (e.g. deposit insurance).

The rest of the paper is organized as follows. In section 2 we review briefly the relevant literature. Section 3 presents the model and it contains the theoretical results for overlapping and random sampling. Section 4 has the simulation results and section 5 concludes.

2 Related Literature

In this section we argue that i) depositors react to the sample of previous decisions that they observe; ii) the law of small numbers describes fairly well how depositors react to what they observe, and iii) correlation across samples matters also. We finish the section discussing briefly the connection to the learning and herding literature.

Most theoretical papers assume that depositors make the decision about withdrawal of funds simultaneously (e.g. Diamond and Dybvig, 1983; Ennis and Keister, 2009). A notable exception is Kinateder and Kiss (2014) who suppose that depositors decide after each other according to a predetermined order (line)

and they observe *all* previous choices before making decision. They study two setups: when, besides actions, liquidity needs of previous depositors are also observed, and when only the previous actions are observed and liquidity needs are private information. In both cases, they obtain that bank runs do not occur in equilibrium.

In this study we suppose that depositors do not observe the decisions of all depositors preceding them, but only a subset. Do people rely on partial information provided by a sample? Kelly and O Grada (2000) test a contagion model which focuses on the banking panics in a New York bank in 1854 and 1857 based on the social network of Irish immigrants. They find that although there were also other factors at work, the most important one determining whether an individual panicked was his county of origin in Ireland. Almost identical persons, only differing in the county of origin, behaved differently during the panics, and the opinions of others with the same background and their choices had the most influence on the decisions of their peers. The origin had its effect through the fact that immigrants from the same county lived in the same neighborhood and observed each other. Iyer and Puri (2012) show that in a bank run episode that took place in India in 2001 a depositors' likelihood of withdrawing is increasing in the share of other people in her social network that have done so earlier. Clearly, a depositor's social network comprises only a reduced subset of all depositors, so only a sample of previous decisions affected if a depositor chose to withdraw or not. Kiss et al. (2014) find in an experiment that observing that somebody has (not) withdrawn increases (decreases) significantly the likelihood of withdrawal.

As already mentioned, the law of small numbers means that agents expect small samples to exhibit large-sample statistical properties, so people tend to overweight information that is available. The main cause behind this phenomenon is that people are too inattentive to the sample size, so they tend to draw overstretched inferences based on small samples (for examples see section 2 in Rabin, 2002). Evidence of people falling prey to the law of small numbers and consequently to overinference abounds when the sample and the population outcome is generated by an exogenous device (e.g. flipping a coin, guessing the urn from which a ball was drawn), but there is less evidence when the observed sample is the outcome of other individuals' decision. However, the findings in Kelly and O Grada (2000) and Iyer and Puri (2012) suggest that depositor behavior can be rationalized by the law of small numbers. More precisely, individuals observing a low number of withdrawals among their peers in their sample may believe that overall the number of withdrawals will be low, so there is no need to run the bank. By the same token, observing many withdrawals in a sample may cause individuals to think that there is a bank run underway. We do not claim that this is the only or the best explanation for their behavior, but it seems a reasonable one. This suggests that a threshold rule may reflect the law of small numbers when studying depositor behavior. In a bank-run

setup, Garratt and Keister (2009) explore experimentally the coordination problem among depositors and find that simple cutoff rules explain fairly well the behavior of the participants.

The arguments in Kelly and O Grada (2000) suggest that correlation of observed behavior is important as people in the same neighborhood had highly overlapping information and they acted also in very similar way. The social network approach in Iyer and Puri (2012) also evokes the relevance of correlated samples as one of the key empirical regularities found in social networks is clustering, that refers to the tendency of linked nodes to have common neighbor(s).⁵ That is, a depositor observes what her neighbors do and at the same time those neighbors are likely to observe each other. Ziebarth (2013) studies the effect of radio penetration on bank runs and banking distress during the Great Depression. He finds that a 10-percentage point increase in radio penetration in a county resulted in a 4.4 percentage point fall in deposits. In our view, radio makes information more correlated, therefore this finding also suggests that the higher the correlation in information, the more likely are massive withdrawals and bank runs.

Regarding learning about the share of previous choices in the population, Costain (2007) is a closely related paper. In his model, individuals choose in a one-shot game between binary actions: playing stock or bond. He assumes that an increase in the fraction of the individuals playing stock increases the payoff related to that choice. Payoffs also depend on the aggregate state of the world and on idiosyncratic shocks. He shows - among others - that if only a sample of previous actions is observed, then multiple outcomes may arise and herds may form. Costain (2007) studies only the case of random sampling, he does not deal with the issue of correlated samples.

Note that this study is not about social learning in which people want to learn the quality of a product or in our case the bank. The model makes it clear in the next section, that the bank is known to work properly, there is no fundamental uncertainty. Thus, if subsequent depositors withdraw in masses, then it is not the consequence of inferring that the bank is bad. In this sense, we do not have a herding model or a global game in which depositors receive noisy private signals about the quality of the bank.⁶ The study rather can be seen as a special coordination game with two types of players in which the players observe a set of previous decisions and decide sequentially.

⁵For more details on clustering in social networks, see for instance Jackson and Rogers (2007) or Watts (1999).

⁶Gu (2011) studies herding in financial intermediation and she assumes that depositors observe some previous actions. Her focus is on the signal extraction problem of depositors about the bank's assets and she assumes away bank runs resulting from coordination failure. Depositors use the information about previous withdrawals to infer the quality of the assets. If their belief about this quality is low, then they withdraw. Since more withdrawals suggest that previous depositors inferred that the assets are not performing well, more withdrawals are more likely to get a depositor to withdraw as well. The idea that observing more withdrawals may start a withdrawal wave is common in Gu's and our paper. However, in our paper the fundamentals are constant and we focus on the possibility of coordination failure.

3 The model

We present a framework based on the canonical Diamond-Dybvig model (1983) with two types of depositors to which we add sequential decision making and sampling. The optimal contract in this setting is such that if only those depositors withdraw who really need liquidity, then those who leave their funds deposited receive more money than those who withdraw. However, the more depositors withdraw, the less money the bank has for future payments. Depositors form a belief about the total number of withdrawals based on what they observe in their sample of previous choices following the law of small numbers and then decide whether to withdraw or keep the money in the bank.

3.1 A modified Diamond-Dybvig model

There is an infinite sequence of depositors who form a bank and deposit their unit endowment there at $T = 0$. The bank invests the deposits in a safe technology which pays unit gross return after each endowment liquidated at $T = 1$ and $R > 1$ after each endowment liquidated at $T = 2$. The long-term return, R , is constant. Therefore, the bank is fundamentally in good conditions and there is no uncertainty in this regard. At the beginning of $T = 1$ a share $0 < \pi < 1$ of the depositors is hit by a liquidity shock and becomes impatient, valuing only consumption in period 1. The rest is of the patient type who enjoy consumption in period 1 and 2. Preference types are not publicly observable and there is no aggregate uncertainty regarding the liquidity preference of the depositors.

At period 0 depositors form a bank by pooling their resources, so that they can share the risk of becoming impatient. Intuitively, the bank will liquidate some long-term investment in period 1 to pay a relatively high (that is, higher than their endowment) consumption to depositors who turn out to be impatient, while patient depositors may enjoy an even higher consumption in period 2. The bank offers this liquidity insurance through a simple deposit contract which is subject to a sequential service constraint. Hence, the bank commits to pay a fixed consumption (c_1) to those who withdraw in period 1 and a contingent payment in period 2 to those who keep deposited their money. The sequential service constraint means that the bank cannot condition the payment on information on the preference type of a depositor who wants to withdraw. Thus, the bank must pay c_1 as a depositor claims her funds back in period 1. The depositors who keep the money deposited receive in period 2 a pro-rata share of the matured assets which have not been withdrawn during period 1.⁷

The timing of events is as follows. At period 0 each depositor deposits in the bank her endowment which

⁷As usual in the literature (Wallace, 1988), in period 1, depositors are isolated and no trade can occur among them.

is invested in the technology. At the beginning of period 1 depositors learn their type privately and nature assigns a position in the line to each depositor. Depositors decide according to this exogenously determined sequence and we suppose that depositors do not know their position in the line and they assign a uniform probability of being at any possible position.⁸ Assume that before decision depositors observe a sample of previous actions which they use to form beliefs about the share of depositors who choose to withdraw. Based on these beliefs, they choose either to withdraw their funds from the bank or keep it deposited (that we also call waiting). We denote withdrawal as action zero ($a_i = 0$) and keeping the money deposited as action 1 ($a_i = 1$).

Denote by (c_1, c_2) the consumption bundle of a depositor in the two periods. Consider the following utility function

$$u(c_1, c_2, \theta_i) = u(c_1 + \theta_i c_2),$$

where θ_i is a binomial random variable with support $\{0, 1\}$. After realization of the needs, if $\theta_i = 0$, then depositor i is impatient caring only about consumption in the period 1, $\theta_i = 1$ stands for a patient depositor. The utility function, $u(\cdot)$ is twice continuously differentiable, increasing, strictly concave, satisfies the Inada conditions and the relative risk-aversion coefficient $-cu''(c)/u'(c) > 1$ for every c .

If types were publicly observable in period 1, then the bank could calculate the optimal allocation based on types and independently of the position in the line. Denote by c_T^θ the consumption of type θ in period T . Then the optimization problem takes the following form:

$$\begin{aligned} & \max \pi u(c_1^0 + c_2^0) + (1 - \pi)u(c_1^1 + c_2^1) \\ & s.t. \\ & (1) \quad c_2^0 = c_1^1 = 0 \\ & (2) \quad \pi c_1^1 + [(1 - \pi)c_2^2/R] = 1 \end{aligned} \tag{1}$$

The first restriction requires that impatient depositors consume in period 1 and patient depositors in period 2. This is optimal, because by consuming in period 2 patient depositors earn the return. The second restriction is the resource constraint. For simplicity, we will denote c_1^{1*} by c_1^* and c_2^{2*} by c_2^* . The solution to this problem is characterized by

$$u'(c_1^*) = Ru'(c_2^*),$$

which implies $R > c_2^* > c_1^* > 1$.

⁸This assumption follows Costain (2007). Introducing (exact or approximate) knowledge about the position would complicate the analysis. Moreover it seems unrealistic to assume that depositors have an accurate view about how many other depositors have already made a decision.

The optimal allocation can be implemented by banks via a simple deposit contract. The depositors withdrawing in period 1 are given c_1^* , while those who keep their funds in the bank reap the benefit of the long-term investment and divide it equally among themselves. The consumption of those who wait depends on the mass of withdrawals in period 1 (ω) and can be expressed in the following way:

$$c_2(\omega) = \begin{cases} \max \left\{ 0, \frac{R(1-\omega c_1^*)}{1-\omega} \right\} & \text{if } 1 - \omega > 0 \\ 0 & \text{if } 1 - \omega = 0 \end{cases}, \quad (2)$$

If $\omega = \pi$, then $c_2(\omega) = c_2^*$. Nevertheless, if ω is high enough, then withdrawing in period 1 is a better option for a patient depositor, than keeping the money deposited (provided the bank has money left to pay). There is a threshold value such that if the number of those who have withdrawn is over this threshold, then the period-2 consumption will be less than c_1^* .

Lemma 1 *There exists a $0 \leq \bar{\omega} < 1$ such that*

$$c_2(\omega) < c_1^* \text{ for any } \bar{\omega} < \omega \quad (3)$$

and

$$c_1^* \leq c_2(\omega) \text{ for any } \bar{\omega} \geq \omega.$$

Proof. The bank cannot pay in period 1 to all depositors $c_1^* > 1$, since depositors have a unit endowment and gross return upon withdrawing in the first period is 1. Hence, for any $\frac{1}{c_1^*} \leq \omega$, $c_2(\omega) = 0$. On the other hand, $c_2^* = c_2(\omega)$ for $\omega = \pi$ and $0 < \frac{\partial c_2(\omega)}{\partial \omega}$ for any $\frac{1}{c_1^*} < \omega < \pi$, so by continuity of $c_2(\omega)$ there is a unique $\bar{\omega}$ such that for any $\omega \leq \bar{\omega}$ we have $c_1^* \leq c_2(\omega)$, whereas for any $\bar{\omega} < \omega$ we have $c_2(\omega) < c_1^*$. ■

Example 1 *Assume that the utility function exhibits constant relative risk aversion (CRRA) and takes the following standard form⁹:*

$$u(c_i) = \frac{c_i^{1-\delta}}{1-\delta}$$

In this case the optimal allocation is the following: $c_1^ = \frac{1}{(1-\pi)R^{\frac{1-\delta}{\delta}} + \pi}$, $c_2^* = \frac{R^{\frac{1}{\delta}}}{(1-\pi)R^{\frac{1-\delta}{\delta}} + \pi}$. The threshold value $\bar{\omega}$ solves the equation $c_1^* = c_2(\omega)$, it is given by the expression $\bar{\omega} = \frac{R - c_1^*}{c_1^*(R-1)}$.*

It is of interest how this threshold varies when the proportion of impatient agents (π), the relative risk aversion coefficient (δ) and the return on the bank's investment (R) change. It is easy to show that $\frac{\partial \bar{\omega}}{\partial \pi} > 0$, hence, the larger is the share of impatient depositors, the higher is the threshold. It is intuitive, because if there are many impatient depositors, then it is more likely that any depositor has a relatively high proportion

⁹This functional form is often used in the bank run literature, see for instance Green and Lin (2003) or Ennis and Keister (2009b).

of withdrawals in her sample. Consequently, patient depositors tolerate more withdrawals in their sample before judging it optimal to withdraw.

The derivative of the threshold with respect to the relative risk aversion coefficient is negative, so the more risk-averse are the depositors, the more they want to smooth the consumption. Consequently, higher risk aversion implies a smaller difference between c_1^* and c_2 , so the threshold decreases.

The results are not clear-cut for $\frac{\partial \bar{c}}{\partial R}$. For different parameter values, the sign of the derivative changes, because an increase in the return has a double effect. On the one hand, it increases the consumption given to those who withdraw in the first period (c_1^*) which - ceteris paribus - lowers the threshold ($\frac{\partial \bar{c}}{\partial c_1^*}$). On the other hand, a higher return increases the period-2 consumption which has a positive effect on the threshold ($\frac{\partial \bar{c}}{\partial R}$). Depending on the parameter values, the first effect either dominates or is dominated by the second one.

3.1.1 Sequential decisions and samples

Our setup differs from Diamond and Dybvig (1983) in two points: depositors decide consecutively according to an exogenously predetermined sequence and before making a decision, they use a sample of previous choices to form beliefs about the share of depositors who withdraw in the first period.

Concerning the size of the sample, for simplicity we focus on the case where each depositor observes a sample of the same size, N .¹⁰ How do the first N depositors decide? They do not have enough choices to observe. For simplicity, we assume that they decide according to their type, that is impatient depositors withdraw, while patient ones keep their money deposited. We consider theoretically two particular cases. One is the random sampling case in which depositor i is equally likely to observe any of her predecessors. The other case has recent predecessors oversampled in a special way. In this case - which we call overlapping sampling - only the last N depositors' actions are observed. This assumption captures, in an extreme format, the idea that recent choices are more likely to be observed. Also, it can be argued that samples may overlap, for example, due to clustering which - as already mentioned - is one of the main stylized facts about social networks.

Depositors use the observed sample to form beliefs about the population share of individuals who withdraw their money from the bank in the first period. To this end, they look at the relative shares of different choices in their sample. Following the behavioral economics literature on the the law of small numbers, we assume that a patient depositor believes that the sample is representative and informative of the whole

¹⁰We suppose that this size is reasonably large. Our results are not driven by excessively small sample size, e.g. observing one or two previous actions. However, the size does not allow to draw a precise conclusion about the share of withdrawals in the whole population.

population.¹¹ For example, if she observes that 60% of her sample withdraws their money from the bank, then she makes some inference based on this information about the share of withdrawals by the end of period 1.

We denote by $\hat{\omega}_i$ the share of withdrawals in depositor i 's sample. To make a decision, depositors compare $\hat{\omega}_i$ to the theoretical threshold value $\bar{\omega}$ defined by Lemma 1. The decision rule can be summarized as:

$$a_i(N, \hat{\omega}_i) = \begin{cases} 1 & \text{if } \hat{\omega}_i \leq \bar{\omega} \\ 0 & \text{if } \hat{\omega}_i > \bar{\omega} \end{cases}, \quad (4)$$

where decision $a_i = 1$ denotes keeping the money deposited, while $a_i = 0$ is withdrawal. If the share of withdrawals in her sample is higher than $\bar{\omega}$, then a patient depositor withdraws. Otherwise, she keeps the money deposited.

A depositor observing a relatively high number of withdrawals believes that what she observes is representative of the proportion of withdrawals at the end of the period. Therefore, it is optimal for her to withdraw her funds from the bank. The rationality of the proposed decision rule may be questioned on the following basis. Our decision rule does not take into account the effect the decision has on the choices of subsequent depositors. This effect is based on the probability that their decision will be sampled by subsequent depositors. Since the samples determine the decision of those depositors, the effect of leaving the money with the bank may be important. The effect is larger for depositors at the beginning of the line and it also depends on the sampling mechanism. Costain (2007) shows in an investment setup that with infinite players inferences about the position are irrelevant for strategies and players can ignore the effects of their own decision on the behavior of others. This lends some support to our modeling choice.

We are interested in whether bank run emerges in our setup or not. A natural way to study this question is to see whether a massive withdrawal wave arises. We define bank run in this paper as a situation in which all depositors withdraw almost surely after some depositor.

Definition 1 *There is a bank run if after some depositor who withdraws all subsequent depositors withdraw almost surely.*

Once there have been depositors who left their money in the bank, it has some probability that a patient depositor happens to observe the decisions of those depositors even though there has been a lot of patient depositors withdrawing in front of her. To account for this possibility, when speaking of a bank run, we

¹¹Impatient depositors always withdraw, so we focus on the decision of patient depositors.

allow for the case that a vanishingly small fraction of patient depositors keep their funds deposited.¹² Note that in the theoretical model with infinite depositors our bank run definition is convenient, since once a bank run starts it implies infinite depositors who withdraw. However, in the case of finite number of depositors a bank run may start late so that the bank's resources will not be depleted. We take care of these issues in the section with the simulation results.

3.2 Overlapping samples

In this subsection, we study the likelihood of bank runs when depositors observe the last decisions. We introduce the following notation. Let $\varphi_i = \hat{\omega}_i N$ and $\tau = \bar{\omega} N$, that is, φ_i denotes the *number* of depositors who withdraw in i 's sample and τ the threshold number of withdrawals in the sample that makes i to withdraw.

First, let us consider how the number of withdrawals evolves with overlapping samples. The number of withdrawals in subsequent samples is

$$\varphi_i = \varphi_{i-1} + (a_{i-(N+1)} - a_{i-1}). \quad (5)$$

where $a_{i-(N+1)}$ is the choice of depositor $i - (N + 1)$. The formula just says that the change in the number of withdrawals in subsequent samples depends on the decisions which are different in the samples. These are the decisions of depositor $i - (N + 1)$ and depositor $i - 1$. The first one is in the sample of depositor $i - 1$, but not in the sample of depositor i , and with the latter one it is the other way around. Hence, the number of withdrawals can differ at most in one unit.

Given any possible threshold τ , it is easy to see that if there has not been any withdrawing patient depositor yet, then the only event that makes a patient depositor to withdraw is to have in her sample more than τ impatient depositors. If $\varphi_{i-1} \geq \tau + 1$, then according to the threshold rule $a_{i-1} = 0$, so depositor $i - 1$ withdraws. Notice that if $a_{i-1} = 0$, then $\varphi_{i-1} \leq \varphi_i$, because $0 \leq a_{i-(N+1)} - a_{i-1}$. This, in turn, implies that $\varphi_i \geq \tau + 1$, so depositor i withdraws as well. By continuing along these lines, we find that once there is a sample such that $\varphi_i \geq \tau + 1$, all subsequent depositors will observe a sample with a number of withdrawals which makes patient depositors to withdraw. Since impatient depositors withdraw always, a sample with more than τ withdrawals sets off a bank run.

Lemma 2 *Given the threshold strategy, if $\varphi_i \geq \tau + 1$ for depositor i , then for any depositor $i \leq j$ we have that $\varphi_j \geq \tau + 1$.*

¹²We could formulate this idea mathematically using a notion of convergence as almost sure convergence. However, it is not necessary to understand the results, so we chose not to include such a formal definition.

Proof. If $\varphi_i \geq \tau + 1$, then $a_i = 0$, so $\varphi_{i+1} \geq \varphi_i$. If depositor i is impatient, then she withdraws, as does a patient depositor i , since $\varphi_{i+1} \geq \varphi_i \geq \tau + 1$. By the same arguments, any subsequent depositor will observe a number of withdrawal higher than τ . ■

Let us consider what happens if we add to this process the proposed decision rule.

Corollary 1 *A bank run starts with the first patient depositor who has at least $\tau + 1$ impatient depositors in her sample.*

Proof. It follows from the definition of bank run and the previous lemma. ■

If no patient depositor has withdrawn yet, then the number of withdrawals in the last N observations follows a binomial distribution with parameters N and π . The lower bound on the probability that a bank run starts within the next N depositors is given by

$$\sum_{x=\tau+1}^N \binom{N}{x} \pi^x (1-\pi)^{N-x} > 0, \quad (6)$$

that is the probability that in the sample there are at least $\tau + 1$ impatient depositors.

Proposition 1 *With overlapping samples and the proposed decision rule, the probability of bank run is 1 and it is independent of τ .*

Proof. The proposition can be proven using the second Borel-Cantelli lemma which asserts that if the events E_i are independent and the sum of the probabilities of the E_i diverges to infinity, then the probability that infinitely many of them occur is 1.

Decompose the sequence of depositors into disjoint blocks consisting of N subsequent depositors. Denote by E_κ the event that at least $\tau + 1$ impatient depositors appear in the κ^{th} block. Since types are distributed independently across depositors and blocks are disjoint, events E_κ are also independent. The probability to have at least $\tau + 1$ impatient depositors in block κ is

$$\Pr(E_\kappa) = \sum_{x=\tau+1}^N \binom{N}{x} \pi^x (1-\pi)^{N-x} > 0. \quad (7)$$

Then $\sum_{\kappa=1}^{\infty} \Pr(E_\kappa) = \infty$, so the probability of having this event infinitely many times is one. For our purposes it is enough to know that it happens at least once and afterwards according to the previous corollary all depositors will withdraw. ■

Discussion

Notice that the results do not change even if patient depositors adjust for the possibility of observing just a sample with too many impatient depositors. The probability of such an event is given by $\sum_{x=\tau+1}^N \binom{N}{x} \pi^x (1-\pi)^{N-x}$. This is the *p-value* of having at least $\tau + 1$ impatient depositors in the sample. Depositors could adjust the threshold (that is, they can choose a higher τ) in such a way that this *p-value* be made very small. It means that depositors want to minimize the error of withdrawing when no run is underway.¹³ The intuition behind this idea is that if somebody observes a sample whose *p-value* is very low, then she has good reasons to believe that it is due to the fact that patient depositors have been withdrawing from the bank. It suggests that a bank run is underway, so the best she can do is to withdraw as well. This kind of adjustment does not help, because as long as $\tau < N$, in an infinite sequence almost surely there will be a sample consisting of only impatient depositors. As a consequence, a bank run starts.

Consider a patient depositor who observes a sample with $\tau < \varphi_i < N$ withdrawals. When observing that at least one depositor has left the money in the bank, the depositor knows that there was a patient depositor before her who - based on her sample - decided to not withdraw. It may suggest that the subsequent depositors who have withdrawn, did it due to being impatient. These considerations are in line with increasing the threshold. In the extreme case, the decision rule may prescribe to keep the money deposited if there is at least one depositor in the sample who did so. We have seen that our result is independent of the threshold, so these arguments do not change the conclusions. Notice also that given an infinite sequence of depositors increasing the sample size does not change the conclusions either.

Note that Proposition 2 extends to the case when depositors differ in their threshold value. This may be the case, for example, when depositors differ in the degree of risk aversion as the threshold $\bar{\omega}$ (and τ) depends on the parameter of relative risk aversion (see Example 1). Suppose that the threshold τ_i is distributed on the interval $[0, N)$ such that the maximum value is $\tau_{max} < N$. In this case a bank run starts when an agent with threshold τ_{max} has $\tau_{max} + 1$ impatient depositors in her sample. This event has positive probability and therefore surely happens in an infinite sequence. All subsequent patient depositors will also withdraw since their threshold is less or equal to τ_{max} .

The result is discouraging, because it tells that bank runs resulting from coordination failures when the last decisions are observable are pervasive. We predict a unique outcome which is just the consequence of chance that determines the type vector according to which depositors decide. There are several arguments

¹³The flipside of the argument is that it may increase the error of not withdrawing when a run is underway. Hence, the optimal extent of adjustment should also take into account the probability of not withdrawing when a run is underway and the loss of doing so.

which palliate this bleak result. The assumption of countably infinite individuals is instrumental to obtain that a bank run arises with certainty. If we consider a model with a finite number of depositors, then the result may change, depending on the size of the population. We run simulations with a finite population of depositors and change the population size between 10^3 and 10^7 . Following Example 1, we use the CRRA utility function and endogenously compute the threshold $\bar{\omega}$. We set the sample size (N) to 25, the second-period investment return (R) to 1.25, the relative risk aversion parameter (δ) to 2.¹⁴ We change the share of impatient depositors between 0 and 1. We run 200 simulations for each parameter setting. Table 1 shows the frequency of bank runs in these 200 simulations. We can see that the probability of a bank run is significantly below 1 when the population is smaller. However, a population of 10^6 depositors very well approximates the case of an infinite population, bank runs occur almost surely.

[Table 1 about here]

In a somewhat similar vein, Temzelides (1997) studies myopic best response in an evolutionary banking setup. In his model with local interaction and without noise depositors observe the share of banks that suffered a bank run in the previous period (and not a sample of previous decisions). He finds that once a bank is run, then possibly a panic ensues involving all banks experiencing a bank run. This is analogous to our result that once a patient depositor withdraws, all subsequent depositors withdraw. However, bank run is not necessary theoretically even if what subsequent depositors observe are similar. Kinaterder and Kiss (2013) show that if depositors decide sequentially and each of them observes all previous decisions, then bank run is not an equilibrium outcome. Therefore, it is clear that it is not correlation in the observed previous choices *per se* that leads to Proposition 1, but correlation and observing only a sample.

3.3 Random samples

In this section, we study the case of random samples. That is, each depositor observes any of the previous decisions with equal probability (and without replacement). If the share of depositors who decided to keep their money deposited up to $i - 1$ is k_{i-1} , then the probability that depositor i will observe a sample of size N with exactly φ_i withdrawals is given by the hypergeometric distribution:

$$\bar{\chi}_i(\varphi_i \mid i, N, k_{i-1}) = \frac{\binom{(1-k_{i-1})(i-1)}{\varphi_i} \binom{k_{i-1}(i-1)}{N-\varphi_i}}{\binom{i-1}{N}}. \quad (8)$$

When the population size is very large, the hypergeometric distribution can be approximated by the

¹⁴This result does not change qualitatively if the sample size is varied. The relative risk aversion parameter has been estimated in many settings and in many countries. Estimates vary substantially, most of the results being between 0 and 4 (see for instance Gándelman and Hernández-Murillo (2014) or Hartley et al. (2013) and the references therein).

binomial distribution:

$$\chi_i(\varphi_i | i, N, k_{i-1}) = \binom{N}{\varphi_i} (1 - k_{i-1})^{\varphi_i} (k_{i-1})^{N-\varphi_i} \quad (9)$$

We study the dynamics of the decisions if depositors follow the threshold decision rule. A patient depositor leaves the money in the bank when observing a sample with at most τ withdrawals, this happens with probability $\sum_{\varphi_i=0}^{\tau} \chi_i(\varphi_i | i, N, k_{i-1})$.

A bank run starts if a patient depositor withdraws and after her only a vanishingly small proportion of depositors keeps their funds deposited. Alternatively, there is a run if after patient depositor i who withdraws any patient depositor j such that $i < j$ follows suit, so the probability of keeping the money in the bank tends to zero. The key issue is to find out whether k_i converges and if it is the case, then to which value. Suppose that for a given threshold decision rule characterized by the threshold (τ) k_i converges to k . As a consequence, $\chi_i(\varphi_i | i, N, k_{i-1})$ converges to $\chi(\varphi | N, k)$.

Given the share of depositors who did not withdraw (k), we define the probability of observing a sample with a number of withdrawals over the threshold as

$$e(\tau, k) = \sum_{\varphi=\tau+1}^N \chi(\varphi | N, k). \quad (10)$$

Note that $0 < e(\tau, k) < 1$ for any $\tau \in (0, N)$. Consider depositor i and suppose that sufficiently many depositors have already decided. The share of depositors who decided to keep their money deposited up to her is k_{i-1} . Then, the share of withdrawals is $(1 - k_{i-1})$ and it equals the share of impatient depositors and the share of those patient depositors who happen to observe more than τ withdrawals. Formally,

$$1 - k_{i-1} = \pi + (1 - \pi)e(\tau, k_{i-1}). \quad (11)$$

As depositor i decides, the share of those depositors who did not withdraw changes to k_i and after each new decision this share varies again. The expression has a recursive structure, and after sufficient decisions the share of withdrawals converges to some $1 - k$. If convergence occurs, then the following condition is met:

$$1 - k = \pi + (1 - \pi)e(\tau, k).$$

After some straightforward manipulation we get

$$k = (1 - \pi) \sum_{\varphi=0}^{\tau} \frac{\binom{(1-k)(i-1)}{\varphi} \binom{k(i-1)}{N-\varphi}}{\binom{i-1}{N}}. \quad (12)$$

This condition means that the share of depositors who keep their money deposited is equal to the share of patient depositors who observe less withdrawals than the threshold τ . Finding k is a fixed point problem.

We are interested in stable crossings, that is values of k where the function represented by the right-hand side in equation (12) crosses the 45° line from above. Notice that $k = 0$ is always a solution, because if the share of depositors who do not withdraw is zero, then all depositors withdraw independently of the threshold, so the proportion of depositors keeping their funds deposited will continue being equal to zero. The question is whether there is a $k > 0$, implying that not all depositors withdraw. There is no bank run if and only if $\exists k > 0$ such that satisfies equation (12).

Given the complexity of $e(\tau, k)$, we could not find conditions that characterize in general when a bank run arises. Therefore, we limit ourselves to numerical methods to assess how large the parameter space is for which bank runs do not occur. We set the sample size (N) to 25 and the return in the second period (R) to 1.25. We change the proportion of impatient depositor π between 0 and 1 by steps of 0.02 and the parameter of relative risk aversion δ between 1 and 4 by steps of 0.02. This gives us 7350 parameter settings. For each we compute the threshold $\bar{\omega}$ using the equations in Example 1 and we check whether there is a positive stable solution $k > 0$ of equation (12). We obtain that in 58.9% of the cases there was no bank run while in 41.1% of the cases there was. Considering random samples thus reduces the likelihood of bank runs by roughly three fifth compared to the case of overlapping samples. Figure 1 illustrates the values of π and δ for which bank runs happened. It shows that bank runs typically occur when either the share of impatient depositors or the second-period investment returns are high. These are the parameter settings for which the endogenous threshold ($\bar{\omega}$) is close to the value of the share of impatient depositors. As it is clear from the proof of Lemma 1, it is always true that $\bar{\omega} > \pi$. When the threshold is very large relative to the share of impatient, bank runs do not occur, since patient depositors need to observe too many withdrawals to decide to withdraw their deposits which is unlikely when π is low. Bank runs happen when the threshold value is near the share of impatient depositors.

Claim 1 *Based on numerical methods we find that with random samples and the proposed decision rule, the probability of bank run is less than 1.*

Figure 1 about here

Discussion

Remember that the bank in our model is assumed to be fundamentally healthy, so bank runs are clearly suboptimal. However, they still occur in the random setting as well, but they are considerably less likely to happen. As indicated earlier, in real life good banks also suffer bank runs and in many occasions one of the driver of the bank run is that depositors observe other depositors rushing to the bank. Our results suggest

that there is some room for preventing runs on good banks by making samples less correlated. More precisely, when massive withdrawals are observed by depositors (for instance through a TV broadcast showing long queues in front of a bank), then the bank or the authority responsible for financial stability should make also visible the decisions of depositors who do not withdraw (for example showing other branches of the bank with no queues or reporting on the stability of deposits in the bank).

The stochastic process implied by random samples is very different from that related to overlapping samples. Most importantly, observing a sample which results in a withdrawal due to the threshold decision rule does not entail that subsequent depositors also will observe a sample with a high number of withdrawals. Thus, the correlation across samples which led to bank runs in the overlapping case is considerably reduced. Nevertheless, random sampling *per se* does not eliminate bank runs. Our preferred interpretation of the previous results is that of comparative statics. Hence, we do not stress the exact likelihood of bank runs, but focus on the finding that less correlation leads to less bank runs.

4 Intermediate cases

We considered two polar cases of sampling and found that while highly overlapping samples lead to bank runs with certainty in our setup, with random samples this need not be the case. It is of interest to know what occurs between these two extremes. As we go from highly to less correlated sampling does the likelihood of bank runs change gradually or are there sharp jumps?

To investigate the role of the degree of correlation/randomness across subsequent samples, we change the degree of randomness, denoted by λ , between 0 (overlapping sample) and 1 (random sample). A depositor observes λN randomly drawn previous decisions and $(1 - \lambda)N$ previous decisions in the line. We study by simulations the frequency and average starting point of a bank run for different parameter values. We consider a population of 10000 individuals. The utility function is CRRA with $\delta = 2$ as in Example 1. We set the investment return in the second period to $R = 1.25$. For each parameter setting we run 500 simulations. Following the model presented in section 3.1, we assume that the first 1000 depositors decide according to their types, that is, impatient agents withdraw while patient depositors keep their money in the bank.

In the first simulation, we set the sample size N to be equal to 25. We change the degree of sample randomness λ between 0 and 1 and the share of impatient depositors between 0.5 and 1. For each parameter setting, we compute the threshold value of the decision rule ($\bar{\omega}$) using the solutions from Example 1.¹⁵ The

¹⁵Notice that for some parameter values, λN is not an integer number. In these cases we use the rounded values in the simulations.

results are shown on Figure 2.

Figure 2 about here

We can observe that for a given share of impatient depositors, as we go from highly correlated towards random samples, the likelihood of bank runs decreases. It is also noteworthy that the change in the likelihood is not gradual in most cases. For instance, when the share of impatient depositors is 0.6 and the sample randomness changes from 0.3 to 0.5, the probability of bank run sharply decreases from 0.83 to 0.02.¹⁶

For the same parameter setting we also calculated the average position in the line where the bank run starts (after which everybody withdraws). Figure 3 shows the results.

Figure 3 about here

In general we can observe that as sampling becomes more random, bank runs start later for any given share of impatient depositors. Transition from an early start to no bank run is quick in some cases. This is desirable since if a bank run begins later, it involves less withdrawals and has less devastating effects. Also, the later the bank run starts, the more time there is to do something to prevent it. Hence, making samples more random has a double effect. On the one hand, it makes bank runs less likely to arise and on the other hand, if there is a bank run, it starts later. These two features together make random samples an important tool to prevent fundamentally unjustified bank runs.

Result 1 *When the sample correlation decreases, bank runs occur less frequently and when they occur, they start later.*

In the second simulation, we consider the interplay between sampling and the sample size. In our model, a bank run starts if subsequent depositors observe many withdrawals. Thus, the more probable that in samples there are more impatient depositors, the more likely is a bank run. If samples are larger, it is less likely to pick too many impatient depositors, so we expect less bank runs with larger samples. We consider again 10000 depositors and the share of impatient depositors (π) is set to 0.8. We change the sample size between 10 and 300 and the sample randomness (λ) between 0 and 0.5.¹⁷ Other parameters of the model are set as before and the threshold $\bar{\omega}$ is computed as in Example 1. We present the results on Figure 4.

Figure 4 about here

There are two noteworthy insights about the likelihood of bank runs when studying the effect of the sample size. First, for a given degree of sample correlation the likelihood of bank runs decreases as the

¹⁶Notice that for the purely overlapping case ($\lambda = 0$), the probability of bank run is not always 1. Proposition 2 stated that for infinite population and overlapping sample bank run occurs with probability 1. In the simulations, however, the population size is 10.000 which is too low to approximate an infinite population in some cases.

¹⁷For larger values of λ bank runs typically did not occur in our simulations, especially when the sample size surpassed 50. Hence we do not show these results on the figures.

sample size increases. Second, for a given sample size the more random is the sample, the less likely are bank runs which indicates that the previous findings are robust to changes in the sample size.

Figure 5 is analogous to Figure 3, it represents the starting point of bank runs in function of sample correlation and sample size. In the lower range of the sample size ($N < 125$), we observe - as before - that increasing the randomness of the sample causes bank runs to start later for any given sample size. However, this effect is diminishing as the sample size gets large. The impact of sample size on the starting point of cascade is rather complex. For more overlapping samples ($\lambda < 0.2$), the starting point of cascade increases when the sample size becomes larger. For more random samples ($\lambda \geq 0.2$), we can observe a non-monotonic relationship between the starting point and the sample size: as the sample size increases, bank runs first start at earlier positions, when sample size reaches higher levels, bank runs emerge at later positions of the line. The impact of sample size is also diminishing: when the sample size is relatively small ($N < 125$), changing it has a significant effect on the starting point. However, for larger values of the sample size, further changes do not significantly affect the starting point of bank runs. For large sample sizes ($N > 125$) starting points for different degrees of correlation (λ) are entangled, with no clear patterns. It suggests that when the sample size increases, correlation becomes less important for the starting point of cascades.

Figure 5 about here

Result 2 *When the sample size increases, bank runs occur less frequently. The relationship between the sample size and the starting point of a bank run is non-monotonic.*

5 Conclusion

Empirical research suggests that depositors' decision is affected by observed previous decisions. The empirical literature also indicates that generally a subset of previous decisions is observed and both type of actions (keeping the money deposited and withdrawal) can be spotted. As the studies analyzing earlier bank run episodes (Kelly and O Grada, 2000; Starr and Yilmaz, 2007; and Iyer and Puri, 2012) show, these features may lead to bank runs even if the bank does not have fundamental problems. We investigate theoretically and using simulations how the way that the sample is collected influences the likelihood of bank runs when depositors make overinferences based on the observed sample. We find theoretically that when comparing overlapping and random samples, bank runs are less likely in the latter case. These findings are corroborated by simulation results that also reveal the importance of the sample size.

Policy-makers can affect both the size and the randomness of the sample. Our results suggest that by requiring to give ampler information about depositors' decision and attempting to make that information

less correlated, the likelihood of bank runs arising from a coordination failure can be efficiently diminished. However, other factors may be relevant also. Schotter and Yorulmazer (2009) find experimentally that bank runs are less severe when participants have more information about other subjects' choices, but only if the economy is in a good state. Ziebarth (2013) suggests that during the Great Depression authorities could have bought time to fix problems by limiting public information. Hence, it seems that in good times more information is better, while in bad times the opposite is true. However, in both cases less correlated information may be helpful to prevent bank runs that are not fundamentally justified.

6 References

Arifovic, J., Hua Jiang, J., Xu, Y. (2013), "Experimental evidence of bank runs as pure coordination failures", *Journal of Economic Dynamics and Control* 37(12), 2446-2465

Barberis, N., R. Thaler (2003), "A survey of behavioral finance", *Handbook of the Economics of Finance*, 1, 1053-1128.

Costain, J.S. (2007), "A Herding Perspective on Global Games and Multiplicity", *The B.E. Journal of Theoretical Economics (Contributions)* 7,

Davison, L.K., C.D. Ramirez (2014), "Local Banking Panics of the 1920s: Identification and Determinants", *Journal of Monetary Economics* 66, 164-177

De Graeve, F., A. Karas (2014). "Evaluating Theories of Bank Runs with Heterogeneity Restrictions", *Journal of the European Economic Association*.12(4), 969-996

DellaVigna, S. (2009), "Psychology and Economics: Evidence from the Field", *Journal of Economic Literature* 47(2), 315-372.

Diamond, D.W., P.H. Dybvig (1983). "Bank runs, deposit insurance and liquidity", *The Journal of Political Economy* 91, 401-419

Ennis, H., T. Keister (2009a). "Bank Runs and Institutions: The Perils of Interventions", *American Economic Review* 99 (4), 1588-1607.

Ennis, H., T. Keister, (2009b). "Run equilibria in the Green-Lin model of financial intermediation", *Journal of Economic Theory* 144 (5), 1996-2020.

Gáandelman, N., R. Hernández-Murillo (2014). "Risk Aversion at the Country Level", *Federal Reserve Bank of St. Louis Working Paper Series (2014-005)*

Garratt, R., T. Keister (2009). "Bank runs as coordination failures: An experimental study." *Journal of Economic Behavior & Organization* 71(2), 300-317.

- Green, E.J., P. Lin, (2003). "Implementing Efficient Allocations in a Model of Financial Intermediation", *Journal of Economic Theory* 109 (1), 1-23.
- Gu, C. (2011). "Herding and bank runs", *Journal of Economic Theory* 146(1), 163-188.
- Hartley, R., G. Lanot, I. Walker (2013). "Who really wants to be a millionaire? Estimates of risk aversion from gameshow data", *Journal of Applied Econometrics*, forthcoming
- Iyer, R., M. Puri (2012), "Understanding Bank Runs: The Importance of Depositor-Bank Relationships and Networks", *American Economic Review* 102(4), 1414-1445.
- Jackson, M. O., B. W. Rogers. (2007) "Meeting strangers and friends of friends: How random are social networks?" *American Economic Review*, 890-915.
- Kahneman, D., A. Tversky (1974), "Subjective probability: A judgment of representativeness", In *The Concept of Probability in Psychological Experiments* (pp. 25-48). Springer Netherlands.
- Kelly, M., C. O Grada (2000), "Market Contagion: Evidence from the Panics of 1854 and 1857", *American Economic Review* 90(5), 1110-1124
- Kinateder, M., H.J. Kiss (2013), "Sequential Decisions in the Diamond-Dybvig Banking Model", mimeo. Available at SSRN: <http://ssrn.com/abstract=2346870> or <http://dx.doi.org/10.2139/ssrn.2346870>
- Kiss, H. J., Rodriguez-Lara, I., Rosa-Garcia, A. (2014), "Do Social Networks Prevent or Promote Bank Runs?", *Journal of Economic Behavior and Organization* 101, 87-99.
- Rabin, M. (2002), "Inference by Believers in the Law of Small Numbers", *The Quarterly Journal of Economics* 117(3), 775-816
- Schotter, A., T. Yorulmazer (2009) "On the dynamics and severity of bank runs: An experimental study." *Journal of Financial Intermediation* 18(2), 217-241.
- Starr, M.A., R. Yilmaz (2007), "Bank Runs in Emerging-Market Economies: Evidence from Turkey's Special Finance Houses", *Southern Economic Journal* 73, 1112-1132
- Sultanum, B. (2014), "Optimal Diamond-Dybvig mechanism in large economies with aggregate uncertainty", *Journal of Economic Dynamics and Control*, <http://dx.doi.org/10.1016/j.jedc.2013.12.011>
- Temzelides, T. (1997), "Evolution, coordination, and banking panics", *Journal of Monetary Economics* 40(1), 163-183.
- Tversky, A., D. Kahneman (1973), "Availability: A heuristic for judging frequency and probability", *Cognitive psychology* 5(2), 207-232.
- Wallace, N., (1988), "Another attempt to explain an illiquid banking system: The Diamond and Dybvig model with sequential service taken seriously", *Federal Reserve Bank of Minneapolis Quarterly Review* 12(4), 3-16.

Watts, D. J. (1999), Small worlds: the dynamics of networks between order and randomness. Princeton University Press

Ziebarth, Nicolas L (2013), "The radio and bank distress in the great depression", mimeo

Figures

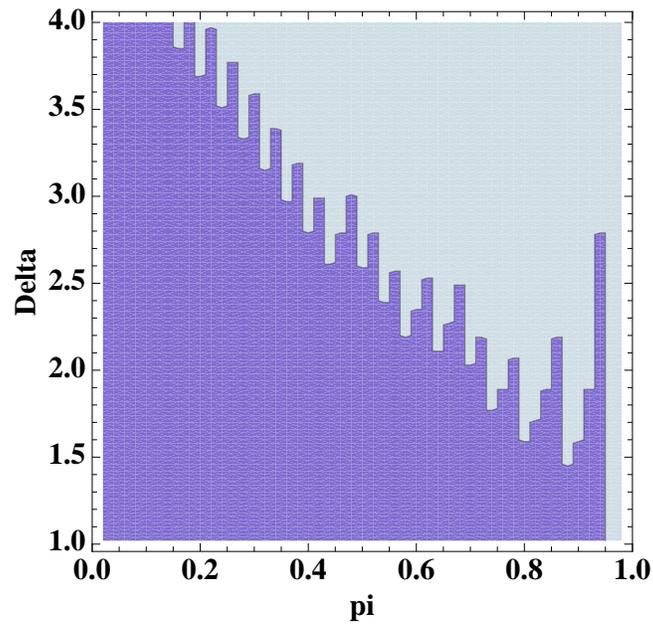


Figure 1: The relationship between parameter values and the emergence of bank runs in the random sampling case for $N = 25$. The light territory shows the cases when bank runs occurred, the dark when they did not occur.

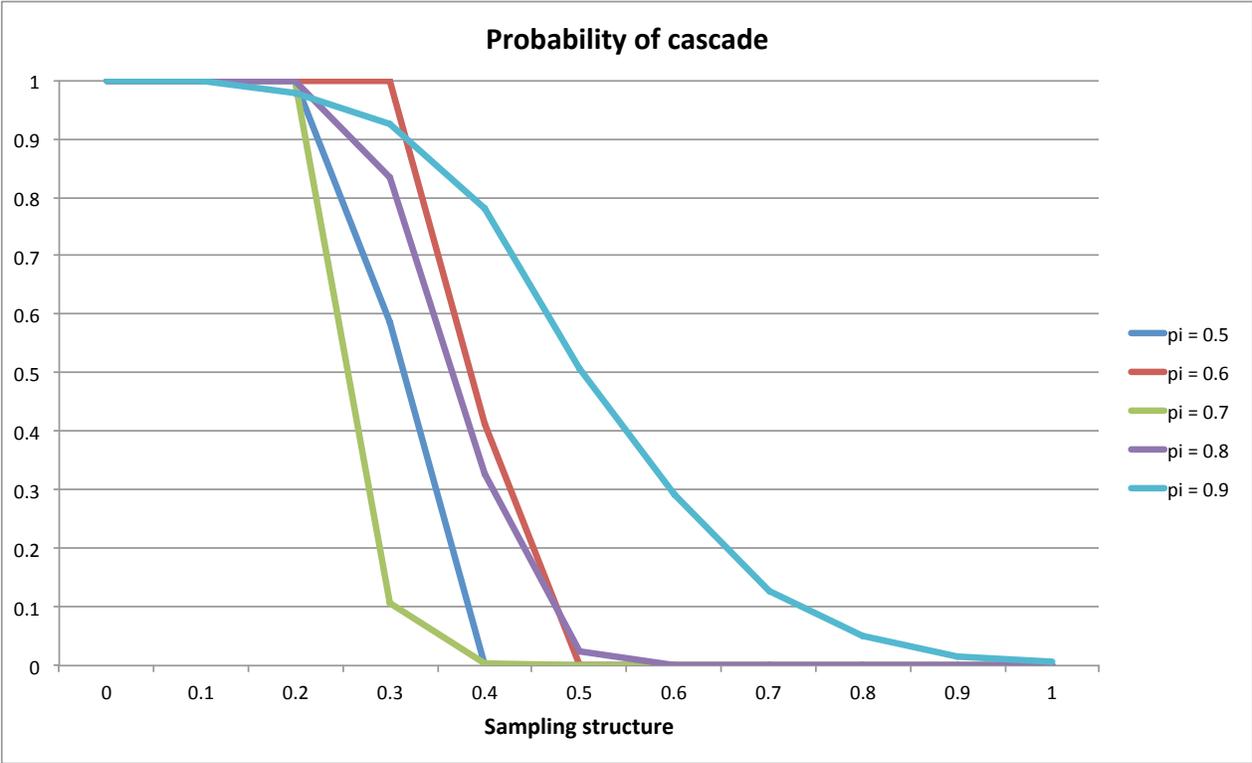


Figure 2: Probability of cascade for different values of sampling parameter λ (x-axis) and share of impatient depositors π .

Tables

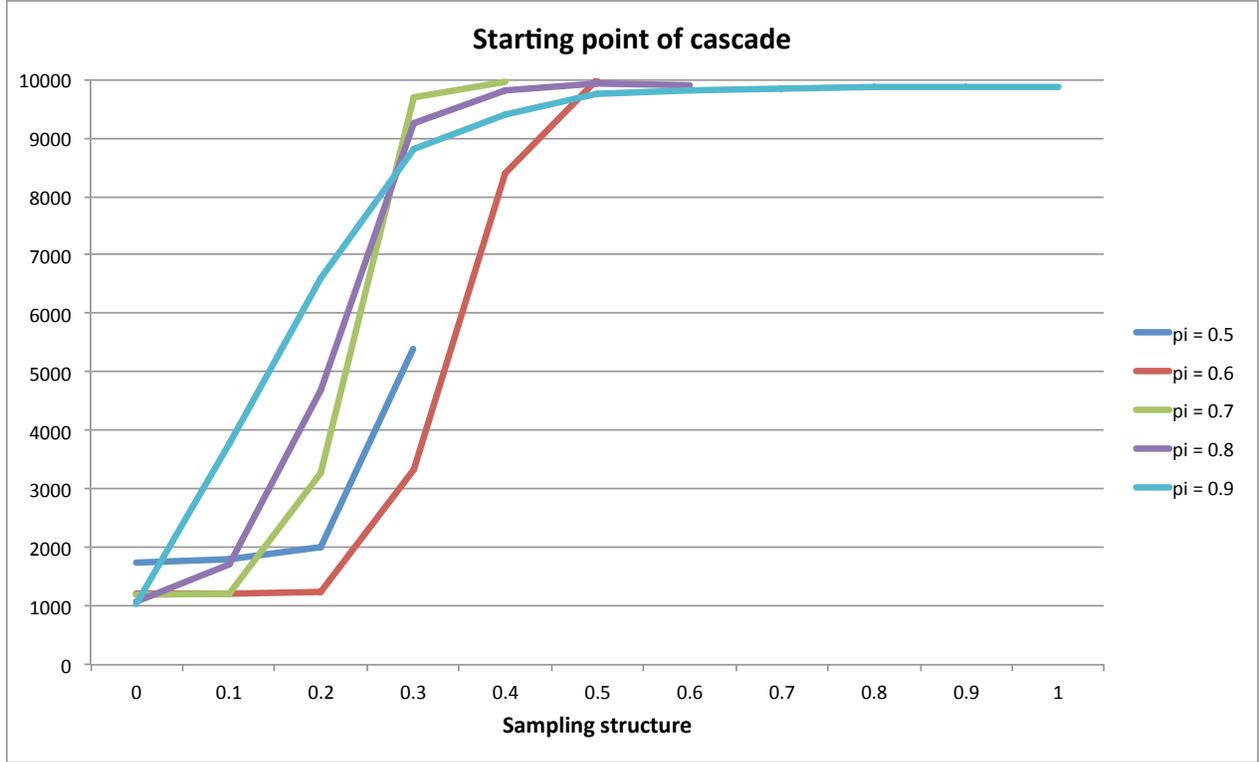


Figure 3: Starting point of cascades for different values of sampling parameter λ (x-axis) and share of impatient depositors π . Average is taken over the simulation runs where bank run emerged.

π	$\bar{\omega}$	Population size				
		10^3	10^4	10^5	10^6	10^7
0.2	0.577	0.015	0.08	0.36	0.995	1
0.4	0.683	0.205	0.94	1	1	1
0.6	0.789	0.985	1	1	1	1
0.8	0.894	1	1	1	1	1

Table 1: Probability of a bank run for finite population size.

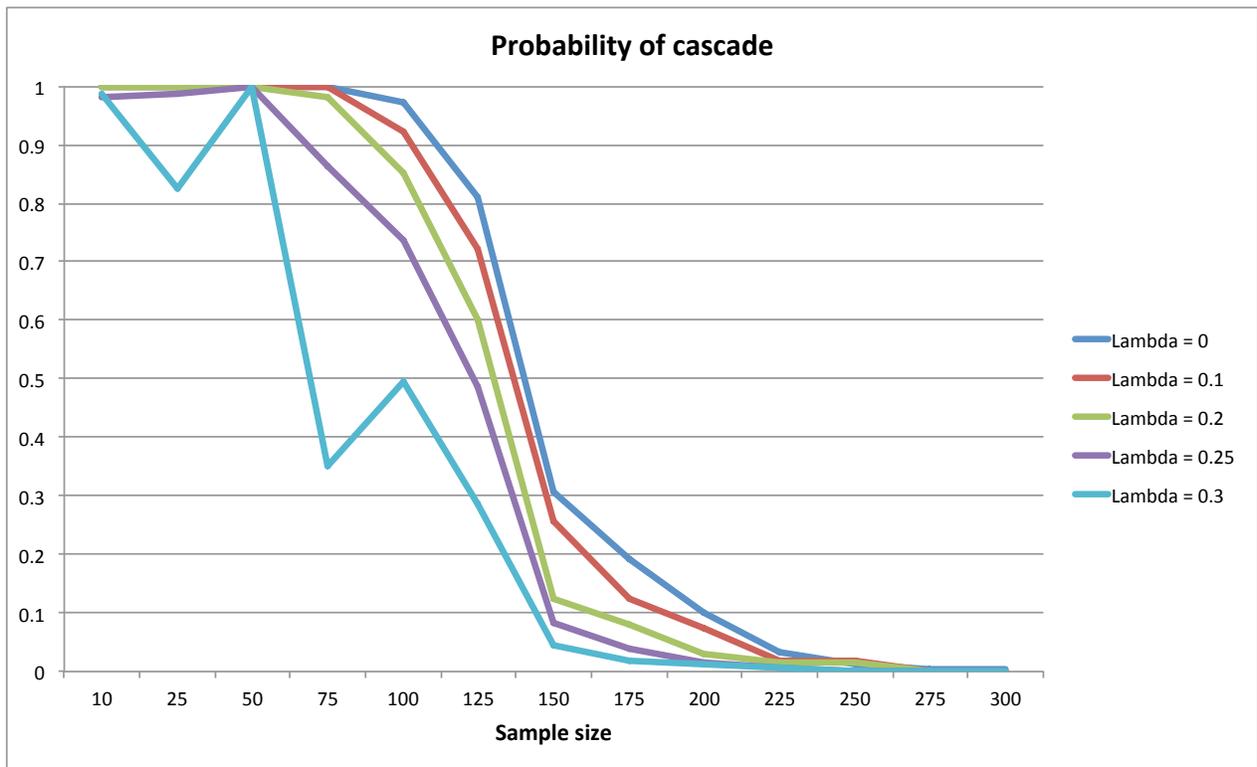


Figure 4: Probability of cascade for different values of sample size N (x-axis) and sampling parameter λ .

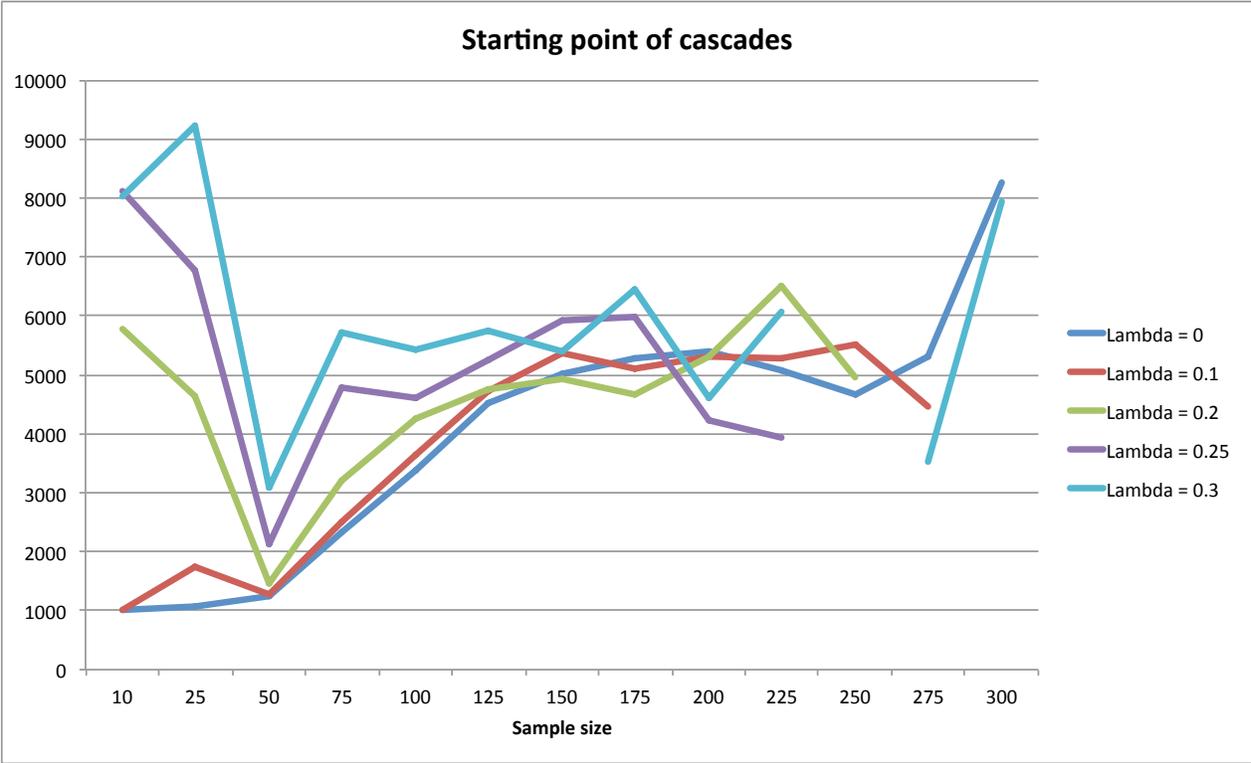


Figure 5: Starting point of cascades for different values of sample size N (x-axis) and sampling parameter λ . Average is taken over the simulation runs where bank run emerged.