Journal of Behavioral and Experimental Finance 44 (2024) 101000

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Finance

journal homepage: www.elsevier.com/locate/jbef

Full length article

Coordination failure in experimental banks of different sizes[☆]

Federico Belotti [a], Eloisa Campioni [a], Vittorio Larocca [b], Francesca Marazzi [c,d], Luca Panaccione [f,*], Andrea Piano Mortari [a,e]

[a] Department of Economics and Finance and CEIS, Tor Vergata University of Rome, Italy
[b] Ministry of Economy and Finance, Department of the Treasury, Italy
[c] Department of Economics and Finance, Tor Vergata University of Rome, Italy
[d] CESifo, Germany
[e] Ministry of Health, Department of Programming, Italy
[f] Department of Economics and Law, Sapienza University of Rome, Italy

ARTICLE INFO

JEL classification:
C72
C92
D80
G21.

Keywords:
Coordination game
Laboratory experiment
Bank run

ABSTRACT

We run a laboratory experiment to investigate how group size affects coordination in a bank-run game, in which participants choose simultaneously whether to withdraw or not and group members change over time. We find that bank size significantly affects the individual withdrawal probability, which is on average 12% higher in large than in small banks. In the initial round(s), all groups exhibit a similar withdrawal rate of about 40%; then, large and medium banks converge to the bank-run equilibrium, while small banks exhibit no systematic convergence. In all banks, experience and beliefs significantly affect the probability to withdraw and to experiment, i.e., to take in the current round the decision opposite to what was the best response in the previous one. We show that experimentation is a strategic choice, and interpret it as an attempt at promoting group convergence towards the efficient equilibrium.

1. Introduction

Runs on banks' deposits can damage economic systems by generating substantial losses and threatening the stability of financial intermediaries.[1] A sound financial intermediary may be vulnerable to a run if depositors coordinate their beliefs and decisions on such an inefficient outcome (Diamond and Dybvig, 1983).[2] What favors or prevents this inefficient coordination has been the object of several theoretical and experimental studies.[3] Indeed, experimental evidence of bank runs is robust to alternative protocols featuring, for example, aggregate uncertainty and multiple withdrawal opportunities (Garratt and Keister, 2009), suspension of deposit convertibility (Madiès, 2006), deposit insurance (Madiès, 2006; Schotter and Yorulmazer, 2009; Peia and Vranceanu, 2019), or observability of fellow depositors' decisions (Kiss et al., 2012, 2014, 2016, 2018, 2022).

A typical feature of the experimental analyses of bank runs is that the dimension of the bank, i.e., the number of participants interacting as depositors, is held fixed across treatment variations. Arifovic et al. (2023) have recently considered whether and how the bank size affects

☆ We thank the associate editor and two anonymous referees for extremely valuable feedbacks. We also thank seminar participants at the IMPRS Workshop in 2018, at the John Cabot University Workshop in 2019, at the Sassari seminar series in 2022, and conference participants at the Experimental Finance Conferences in 2017, 2018 and 2021 and at the FUR 2018. This research has benefited from financial support of the MIUR (PRIN 2015 and PRIN 2022). The paper supersedes a previous version titled "Born to Run: Adaptive and Strategic Behavior in Experimental Bank-Run Games" The views expressed in this paper are those of the authors and should not be attributed to the Italian Ministry of Economy and Finance.

* Corresponding author.
E-mail addresses: federico.belotti@uniroma2.it (F. Belotti), eloisa.campioni@uniroma2.it (E. Campioni), vittorio.larocca@mef.gov.it (V. Larocca), francesca.marazzi@uniroma2.eu (F. Marazzi), luca.panaccione@uniroma1.it (L. Panaccione), piano.mortari@economia.uniroma2.it (A. Piano Mortari).

[1] In the last decades, runs occurred in Brazil (1990), Russia (1995), Malaysia (1999), Ecuador (1999), Argentina (2001) and Uruguay (2002), UK (2007), United States (2008) among others. In the US, in March 2023 the Silicon Valley Bank collapsed due to a bank run involving an outflow of deposits of about 40 billion US dollars (Barr, 2023).

[2] An alternative perspective has emphasized that runs on banks may also be associated to weak fundamentals (Saunders and Wilson, 1996; Calomiris and Mason, 2003; Allen and Gale, 1998), but this is not the case in the experiment we designed.

[3] For the theoretical analyses, see Green and Lin (2003), Peck and Shell (2003), Goldstein and Pauzner (2005), Allen and Gale (2009), Ennis and Keister (2009), Andolfatto et al. (2017) and Dávila and Goldstein (2023) among others. For the experimental analyses, see Dufwenberg (2015) and Duffy (2016) for comprehensive surveys.

https://doi.org/10.1016/j.jbef.2024.101000
Received 2 November 2023; Received in revised form 3 October 2024; Accepted 23 October 2024
Available online 10 November 2024

the coordination of subjects' decisions. In particular, they compare very large (for experimental standards) groups of 75–90 depositors with groups of 10 depositors, and find relevant differences in subjects' behavior across bank sizes. This paper shares the same comparative perspective of Arifovic et al. (2023), and contributes to the experimental analysis of individual decisions in a bank-run game à la (Diamond and Dybvig, 1983).

Within a similar setting, Arifovic et al. (2013) have analyzed the tightness of the coordination problem by varying the coordination parameter, i.e., the share of fellow depositors' withdrawals which makes a depositor indifferent between withdrawing or not. Their results show that, holding the bank size constant, (in)efficient outcomes are more likely when the parameter is (low) high.[4] We use a modified version of the game implemented by Arifovic et al. (2013, 2023), and consider three bank sizes, namely five, seven, and ten depositors. To standardize the tightness of the coordination problem, in all banks we use the same parameter, and we set it to 35%. This value belongs to the interval in which Arifovic et al. (2013) find no systematic convergence to either the bank-run or the no-run equilibrium, and it is close to the one for which Arifovic et al. (2023) find that all very large groups converge to the bank-run equilibrium, while the smaller ones do not exhibit a clear pattern of convergence.

Differently from Arifovic et al. (2013, 2023), who use a partner protocol, in our experiment participants interact repeatedly with possibly different depositors (stranger protocol), and decide simultaneously whether to withdraw or not, being aware of past outcomes. These features mimic relevant characteristics of contemporary digital accounts and multiple-banking, with depositors managing several accounts and interacting anonymously in part with the same fellow depositors over time.[5] In these cases, the experience matured in a bank can affect a depositor's account management in other banks, where one cannot exclude that she interacts with the same fellow depositors. Furthermore, we include in the experiment an explicit belief elicitation task regarding the fellow depositors' choices. This allows us to analyze how decisions are affected by experience and strategic considerations, and whether their effects vary with the size of the bank. To investigate the role of experience in the dynamics of the individual withdrawal choices, we also rely on the analysis of Arifovic et al. (2013),[6] and therefore distinguish between ex-post rational decisions that are best responses to previous-period outcome,[7] and experimentation (in their terminology), i.e., decisions that are opposite to those best responses. Given the aforementioned characteristics, our comparative experimental analysis provides a robustness test of the findings of Arifovic et al. (2023) for smaller bank sizes and for the possibility of having new fellow depositors in one's group.

Our experimental data show that, in early rounds, withdrawal rates are fairly high (40%) and rather homogeneous across bank sizes. In later rounds, the behavior becomes more heterogeneous: in small (five-depositor) banks the rate of withdrawals remains stable over rounds, while in medium (seven-depositor) and large (ten-depositor) ones, it increases steadily up to reaching almost universal withdrawing in the final rounds. This evidence is consistent with Arifovic et al. (2023) in that outcome indeterminacy emerges if the group is small, otherwise there is a clear convergence process toward the inefficient outcome. Furthermore, we find that in all groups subjects tend to correlate their expectations to the observed fellow depositors' choices in previous rounds, and take decisions which are consistent with the formulated expectations. Regarding the role of experience, the majority of participants align their current choices to previous-period withdrawal rate (ex-post rational behavior), while a non-negligible fraction engages in experimentation. Experimentation is mostly realized by choosing not to withdraw, and it counts up to about half of the participants' choices in all groups; subsequently, it remains fairly stable in small banks, while in medium and large ones it declines across rounds.

To perform our econometric analyses, we design a strategy that allows to study the role of bank size while controlling for the cross-sectional correlation generated by the rematching protocol. Specifically, all our models include a proxy for the behavior of fellow depositors in previous round(s), which is the main source of such correlation. Once this is controlled for, the econometric analyses confirm that the bank size significantly affects both the probability to withdraw and to experiment.

The econometric analyses show that the probability to withdraw increases with bank size: in large (medium) banks it is on average 12% (10%) higher than in small ones. Moreover, the withdrawal probability increases with expected withdrawals by fellow depositors and also with observed previous-round withdrawals.

We also find that group size is a determinant of the probability to experiment, and the two exhibit an inverse relationship. Specifically, the experimentation probability is on average 27% lower in large banks relative to small ones. We also find that a participant is more likely to experiment if the previous-period observed withdrawal rate is closer to the coordination parameter, as in Arifovic et al. (2013), and if the expected withdrawals are lower than those realized in the previous period, even after learning across rounds is accounted for. This evidence supports the idea that experimentation is not a random choice, but emerges as a strategic decision of those participants who attempt at promoting group convergence towards the efficient outcome.

The paper is organized as follows: Section 2 presents the experimental game, the main hypotheses, and the protocol. Section 3 presents the econometric strategy. Section 4 discusses the main results. Section 5 concludes. Appendix A contains additional tables and figures. Appendix B contains supplementary material, including the instructions.

2. Experimental game and hypotheses

In our experimental game, adapted from Arifovic et al. (2013), participants ("depositors") are assigned a bank deposit and simultaneously choose to withdraw it ($\boldsymbol{w}$), or not ($\boldsymbol{n}$). The bank invests its resources and promises the return R at maturity, and a lower return r, with $1 \leq r$, in case of premature liquidation. We assume that (i) the payoff from $\boldsymbol{n}$ is greatest when all depositors choose $\boldsymbol{n}$ and decreases with the number of withdrawals; (ii) the payoff from $\boldsymbol{w}$ is constant as long as the bank has resources to repay r. Assumption (i) implies that coordination on $\boldsymbol{n}$ is Pareto efficient and that withdrawals trigger costly premature liquidation of the bank's assets. Assumption (ii) implies that the bank accommodates withdrawals by depleting its resources and that these may be insufficient to repay r if too many choose $\boldsymbol{w}$.

Let N be the total number of depositors and, also, the total value of deposits.[8] Moreover, let N_w and N_n, with $N_w + N_n = N$, be the

[4] This is consistent with the fact that the lower the parameter, the smaller the share of withdrawals which renders withdrawing preferable to not withdrawing, hence the easier the coordination on the bank-run outcome.

[5] Multiple-banking refers to the phenomenon that a depositor holds two or more current accounts in different banks. It is well-documented by researchers in marketing and retail services (see for instance, Gerrard and Cunningham (2001) and Devlin and Gerrard (2005)). It occurs in EU countries, US and Asia since the late 1990s. In the period 2013–2015, in continental Europe the fraction of residents who hold two or more current accounts is above 40% and in UK it is around 30% (see Evans (2015)).

[6] Arifovic et al. (2013), in turn, build on the evolutionary model of Temzelides (1997) who considers a repeated version of the Diamond and Dybvig (1983) game.

[7] See Selten et al. (2005, p.6): "The way in which the decision is based on experience may be described as "ex-post rationality". One looks at what might have been better last time and adjusts the decision in this direction".

[8] This implies that the nominal value of each deposit is equal to 1 (unit of numeraire).

Table 1
Payoff table for banks with $N = 5$.

	Payoff if you withdraw ∘	Payoff if you do not withdraw •
∘∘∘∘	98	7
•∘∘∘	122	90
••∘∘	122	117
•••∘	122	132
••••	122	150

number of depositors who choose $\boldsymbol{w}$ and $\boldsymbol{n}$, respectively. Furthermore, let $\pi_n \equiv \pi(\boldsymbol{n}|(N_n, N_w))$ and $\pi_w \equiv \pi(\boldsymbol{w}|(N_n, N_w))$ be the individual payoff from $\boldsymbol{n}$ and $\boldsymbol{w}$, respectively. Finally, let $0 \leq \lambda$ and $\psi < 1$ be the cost of managing financial investments and the default cost, respectively.[9] Therefore, the depositors' payoffs are given by:

$$\pi_n = \begin{cases} \max\left\{0, \dfrac{R(N - rN_w)}{N_n} - \lambda\right\} & \text{if} \quad 1 \leq N_n < N, \\ R & \text{if} \quad N_n = N; \end{cases} \tag{1}$$

and

$$\pi_w = \begin{cases} \min\left\{r, \dfrac{N}{N_w}\right\} & \text{if} \quad 1 \leq N_w < N, \\ 1 - \psi & \text{if} \quad N_w = N. \end{cases} \tag{2}$$

If $N_w = 0$, the total return RN is shared equally among depositors, yielding an individual net return R. If $1 \leq N_w < N$, the bank must liquidate rN_w to repay r to those who choose $\boldsymbol{w}$. The remaining amount $R(N - rN_w)$ is equally distributed to those who choose $\boldsymbol{n}$, for an individual return, net of cost λ, equal to $(R(N-rN_w)/N_n)-\lambda$. The return r is guaranteed as long as $r \leqslant N/N_w$; otherwise the bank is unable to repay r and, in this case, those who choose $\boldsymbol{w}$ get the fraction N/N_w of the available resources. Finally, when $N_w = N$, all resources are liquidated and shared equally among depositors net of a default cost ψ, so that $\pi_w = 1 - \psi$.

In the experiment, the deposit is equal to 100 units, $R = 1.5, r = 1.22, \lambda = 0.1$, and $\psi = 0.1/N$. Given these values and using Eqs. (1) and (2), Table 1 reports the payoffs when $N = 5$, with white circles denoting $\boldsymbol{w}$ and black circles denoting $\boldsymbol{n}$.

Recall that the coordination parameter q_w is the share of fellow depositors' withdrawals which makes a depositor indifferent between $\boldsymbol{w}$ and $\boldsymbol{n}$. We consider a common (across bank sizes) coordination parameter of $q_w = 0.35$, since this is in the range for which Arifovic et al. (2013) find no systematic coordination of subjects' choices.[10] Therefore, when $N_w/N > 0.35$ the (unique) best response is $\boldsymbol{w}$ (first to third row in Table 1), while when $N_w/N < 0.35$ (fourth and firth row in Table 1) the (unique) best response is $\boldsymbol{n}$. It follows that there are two (pure-strategy) Nash equilibria: one with $N_w = N$, i.e., with a bank run, and payoff equal to 98, and one with $N_w = 0$ and payoff equal to 150.

In this theoretical benchmark, the individual strategic behavior depends only on the expectations about fellow depositors' choices. Even when applying alternative equilibrium refinements, such as those examined by Arifovic et al. (2023, Table 1 in Appendix A), the size of the bank does not affect equilibrium behavior. Therefore, in the absence of theoretical predictions regarding the role of the group size, we refer to the available experimental evidence to formulate our behavioral hypotheses.

Using a similar experimental game, Arifovic et al. (2023) compare withdrawal choices in groups of 10 and of 75–90 depositors. With a coordination parameter comparable to ours ($q_w = 0.30$), they find that participants withdraw more often in large than in small groups, and that all large groups converge to the run equilibrium while results are mixed for small ones. Hence, we formulate the following:

Hypothesis 1. Participants withdraw more often in larger than in smaller groups.

Experimental evidence for bank-run games (see, e.g., Garratt and Keister (2009), Kiss et al. (2014), and Arifovic et al. (2023)) supports the idea that the withdrawal probability increases with the frequency of observed past withdrawals, even with a stranger protocol. Past interactions may affect current decisions in at least two ways: observed outcomes could be reliable predictors of current outcomes, or could induce attempts at changing the outcome towards which subjects seem to converge. We can interpret the first effect as measuring whether participants choose to withdraw because in the previous period it was the optimal decision to take (ex-post rational behavior). Any attempt at changing group behavior would imply that a subject switches choice relative to the previous period best response, hence that experimentation has strategic determinants, in line with what has been first noticed by Arifovic et al. (2013).[11] If this is the case, one would expect that the size of the group affects the likelihood of experimentation, for instance, through the perception of being less pivotal as the group gets larger (see Arifovic et al. (2023)). Moreover, a participant could experiment more in the initial periods when a clear convergence process has not taken off yet. Both considerations are supported by the evidence in Arifovic et al. (2023) that participants in small groups exhibit a more heterogeneous behavior and that in large groups there is greater persistence of the withdrawal choices. We therefore formulate the following:

Hypothesis 2. The probability of experimentation is higher in initial periods and in smaller groups.

2.1. Experimental protocol

The experiment is organized in three phases (Phase 1, Phase 2 and Phase 3).[12] Participants receive feedback information after each phase and earnings are determined at the end of the experiment.

In Phase 1, participants answer a multiple-choice questionnaire (13 questions) on financial literacy and general knowledge.[13] We elicit participants' financial literacy to test whether this individual characteristic affects decisions in the bank-run game, given also that the instructions follow a financial frame. The time limit to answer each question is set to 90 seconds. Wrong answers are penalized and unanswered questions are neither rewarded nor penalized. The total score is converted into the probability of winning the high prize (150 Zed) instead of the low one

[9] In the experiment, the financial cost λ is borne by depositors who do not withdraw, except in the default case, when every depositor incurs a net loss of ψ in the nominal value of its deposits.

[10] Using (1) and (2), and given the values of the parameters, one can verify that $q_w = 0.35$ equals the payoffs from the two decisions.

[11] This notion of experimentation does not prevent the subject's choice to be optimal given the beliefs.

[12] The Instructions are provided in Appendix B.

[13] The questionnaire consists of six questions on financial literacy. Three questions focus on the notions of inflation, of share, and of compound interest rate, and are adapted from the Basic and Advanced Literacy Questions in van Rooij et al. (2011). One question deals with the pricing of an asset, and it is adapted from PISA 2012 Financial Literacy Questions and Answers, proposed by OECD (2012a). The last two questions relate to portfolio decisions and to inter-temporal choices and are proposed in an original formulation. The seven general-knowledge questions are adapted from the PISA released items on mathematics, problem-solving, and field trial cognitive abilities (see OECD (2012c,b, 2015)).

Table 2
Summary of experimental sessions.

Session	Treatment	Number of participants	Rounds	Number of banks per round	Condition
1	5-depositor banks	25	20	5	Table I
2	7-depositor banks	21	20	3	Table I
3	10-depositor banks	20	25	2	Table I
4	5-depositor banks	25	20	5	Table D
5	7-depositor banks	21	20	3	Table D
6	10-depositor banks	20	25	2	Table D

(50 Zed) in a binary lottery.[14] Feedback information includes the score attained for the two groups of questions and the winning probability.

In Phase 2, participants play the bank-run game. Since our treatments consist in varying the size of the experimental banks, in separate sessions we have banks with 5, 7 and 10 depositors respectively, i.e., small, medium and large banks. The bank-run game is repeated for 20 rounds in small and medium banks and for 25 rounds in large banks.[15] Participants interact anonymously and are re-matched in each round (stranger protocol).[16] They have to decide whether to withdraw or not within 30 s, otherwise the program will randomly implement a decision.[17] We elicit non-incentivized expectations about withdrawal choices of fellow depositors, similarly to Kiss et al. (2022), as a preparatory task for the withdrawal decision, without imposing any time limit. End-of-round feedback includes the own payoff and the number of withdrawals in the own bank. One round is randomly selected for payment.

In Phase 3, we elicit participants' risk aversion using the Holt and Laury (2002) protocol. Prizes for the safe lottery are 200 Zed and 160 Zed, while prizes for risky lottery are 385 Zed and 10 Zed, so that the magnitudes are comparable with the payoffs in the bank-run game.

The actual payment is determined after Phase 3, and it is equal to the sum of (i) the prize of the binary lottery associated to Phase 1, (ii) the payoff of the randomly selected round of Phase 2 and (iii) the prize of the binary lottery from the risk-aversion elicitation task. The binary lotteries are played by the computer. The payment round for Phase 2 and the lottery pair for Phase 3 are randomly selected using a public device (bingo numbers). The average payment was 23 euros.

Given the extensive evidence of framing effects (see, for instance Weimann and Brosig-Koch (2019, Ch. 2.5.3)), we implement two between-subject conditions which only differ in the order in which payoffs are displayed in the payoff tables, increasing in condition Table 1 and decreasing in condition Table D (see Tables B.1–B.3 and Tables B.4–B.6 in Appendix B).

The experiment was programmed with z-Tree (Fischbacher, 2007) and conducted in the CESARE lab at LUISS University (Rome, Italy) between November 2015 and March 2016. 132 subjects, recruited with Orsee (Greiner, 2015), participated in only one of the overall six sessions (see the overview in Table 2). The sample of participants is fairly balanced in terms of gender (48.1% are female participants) and the average subject is less than 22 years old, risk averse and performs fairly well both in the financial literacy and the general knowledge questionnaire. In terms of education, 68.3% of students is from Economics, 17.25% from Law, 13.03% from Political Science and 1.4% from other domains.

[14] The winning probability ranges from 5% to 95% as the score increases from −6.5 to 13 points. The denomination of the currency is borrowed from OECD (2014).

[15] Rounds are adjusted to have a similar amount of individual observations in the groups of different size, given the capacity of the lab. This implies, in particular, that the number of individual observations (participants times rounds) is equal in small and large banks.

[16] We do not implement a perfect stranger protocol.

[17] Participants made all their decisions within the given time limit.

Table 3
Experimentation of subject i, $e_{it}(n)$ and $e_{it}(w)$.

		$(t-1)$-*best response*	
		w	n
choice at t	w	$e_{it}(n) = 0$	$e_{it}(w) = 1$
	n	$e_{it}(n) = 1$	$e_{it}(w) = 0$

3. Econometric strategy

To test our hypotheses, we design an econometric strategy which relies on multivariate regressions and includes, in every specification, a proxy for the cross-sectional correlation generated by the matching protocol implemented in the experiment. In particular, we construct variables about group's choices in the previous round that we believe are able to capture such correlation, as it will be explained below. Furthermore, since our setting is characterized by a low number of independent observations but a large number of individual ones, we resort to regression analysis rather than other commonly used approaches, such as parametric tests, to properly analyze the role of bank size on participants' decisions, while controlling for the aforementioned correlations.

Concerning the withdrawal decisions, two main drivers are key to examine: one related to experience, and the other to strategic deliberation based on expectations. Furthermore, current decisions may be affected by own past withdrawal decision (pure state dependence) and by unobserved individual factors (spurious state dependence). If unobserved heterogeneity is correlated over time, own past decisions may appear a determinant of current decisions solely because they are a proxy for such temporally persistent unobservables (Heckman, 1981). To disentangle the effects of pure and spurious state dependence, we use the following dynamic correlated random-effects probit model:

$$Pr(y_{it} = 1 | x_{it}, \boldsymbol{z}_i, y_{it-1}, c_i) = \Phi(\gamma y_{it-1} + x_{it}\beta + \boldsymbol{z}_i\boldsymbol{\theta} + c_i + \tau_p) \quad (3)$$

$$c_i | x_{it}, y_{i1} \sim \mathcal{N}(\alpha + \nu\bar{x}_i + \eta y_{i1}, \sigma_u^2), \quad (4)$$

where $y_{it} \in \{0,1\}$ is the binary variable indicating the withdrawal decision of depositor i in round t. The time-varying covariates are the individual withdrawal decision in the previous round (y_{it-1}) and, depending on the model specification, the variables (x_{it}) controlling either for the share of expected withdrawals from fellow depositors or for the share of fellow depositors who withdrew in previous rounds. We control for time effects by clustering rounds in groups of 5 using τ_p with $p = 1, \dots, 5$.[18] The time-invariant covariates ($\boldsymbol{z}_i$) include dummies for bank size, a set of demographic controls, and a dummy for condition Table D.[19] In our model, Eq. (4) follows from $c_i = \alpha + \nu\bar{x}_i + \eta y_{i1} + u_i$, with $u_i \sim \mathcal{N}(0, \sigma_u^2)$ and independent of y_{i1}, i.e., the first-round withdrawal decision, and of $\bar{x}_i$, i.e., the average of the time-varying covariates at individual level. In this specification, c_i controls for time-invariant unobserved heterogeneity that may be correlated with x_{it} and y_{i1}. As pointed out by Wooldridge (2005), once the initial conditions are controlled for, the parameters of this model can be consistently estimated by conditional maximum simulated likelihood.[20]

[18] In banks with 5 and 7 depositors, $p = 1, \dots, 4$, as the sessions consist of 20 rounds.

[19] The condition dummy, and its interactions with the grouped round dummies, account for the possible effects of the Table conditions based on the evidence from Table A.1 and Fig. A.1 in Appendix A. Specifically, Table A.1 reports the withdrawal rates by condition, both overall and for the first round separately, and, for the latter only, a two-sample proportion test on the between-condition difference. Although this is not significantly different from zero, the data suggest that withdrawals are lower in condition Table D than in Table 1, and that this trend is persistent throughout rounds (see Fig. A.1).

[20] Estimation of (3)–(4) has been performed using the Stata command `xtprobit, re`.

Table 4
List of variables.

Variable	Mean	sd	Description
Withdraw	0.733	0.442	Dummy: subject has decided to withdraw in current round
Expectation_t	0.669	0.333	Share of fellow depositors who are expected to withdraw at t
Feedback_{t-1}	0.728	0.285	Share of fellow depositors who withdrew at $t-1$
N5	0.352	0.478	Dummy: 5-depositor bank treatment
N7	0.296	0.456	Dummy: 7-depositor bank treatment
N10	0.352	0.478	Dummy: 10-depositor bank treatment
Table D	0.500	0.500	Dummy: table with decreasing payoffs
Distance_{t-1}	0.336	0.278	Absolute value of difference between Feedback_{t-1} and q_w
δ^-	0.355	0.479	Dummy: $\text{Expectation}_t < \text{Feedback}_{t-1}$
δ^+	0.272	0.445	Dummy: $\text{Expectation}_t > \text{Feedback}_{t-1}$
Demographics			
Female	0.481	0.500	Dummy: subject is a female
safeCh	6.055	1.780	Number of safe choices made in risk-elicitation task
Age	21.85	2.243	Age
Economics	0.683	0.465	Dummy: subject is a student in Economics
eduMother	0.525	0.499	Dummy: subject's mother has academic-level education
scoreFin	2.627	1.616	Score of the subject on questions in financial literacy
scoreGen	3.104	2.045	Score of the subject on questions in general knowledge

For comparability across bank sizes, $\texttt{Expectation}_t$ is equal to expected, and $\texttt{Feedback}_{t-1}$ to actual withdrawals by fellow depositors divided by the number of fellow depositors, i.e., $N-1$, with $N = 5, 7, 10$. The frequencies of δ^- and δ^+ do not sum up to 1 as the complementary frequency of 0.373 corresponds to cases where $\text{Expectation}_t = \text{Feedback}_{t-1}$.

In line with Arifovic et al. (2013), we identify experimentation in round t with reference to what would have been the optimal choice in the previous round, i.e., the $(t-1)$-best response. Therefore, we say that subject i experiments with not withdrawing if she chooses $\boldsymbol{n}$ at t when withdrawing would have been the $(t-1)$-best response, i.e., if the observed withdrawal rate by fellow depositors is larger than q_w. Similarly, we say that subject i experiments with withdrawing when she chooses $\boldsymbol{w}$ at t if not withdrawing would have been the $(t-1)$-best response. Since we are interested in both types of experimentation, we let $e_{it}(w) = 1$ if subject i in round t experiments with $\boldsymbol{w}$, $e_{it}(n) = 1$ if she experiments with $\boldsymbol{n}$, and $e_{it}(n) = e_{it}(w) = 0$ if she does not experiment in round t (see Table 3).

This definition of experimentation results in two unbalanced panels with gaps, thus precluding a reliable dynamic analysis. Therefore, we consider the following (static) correlated random-effects probit model:

$$Pr(y_{it} = 1|\boldsymbol{x}_{it}, \boldsymbol{z}_i, c_i) = \Phi(\boldsymbol{x}_{it}\beta + \boldsymbol{z}_i\boldsymbol{\gamma} + c_i + \tau_p) \quad (5)$$

$$c_i|\boldsymbol{x}_{it} \sim \mathcal{N}(\alpha + \nu\bar{\boldsymbol{x}}_i, \sigma_u^2), \quad (6)$$

where $y_{it} \in \{e_{it}(n), e_{it}(w)\}$ is the experimentation of depositor i in round t while $\boldsymbol{z}_i$, τ_p and c_i are as in (3)–(4). We include in $\boldsymbol{x}_{it}$ a control capturing the distance between the share of withdrawals observed by subject i in $(t-1)$ and the threshold parameter q_w, and a dummy that compares expectations in round t with the observed withdrawals in $(t-1)$. Specifically, when estimating the probability of experimenting with $\boldsymbol{w}$, the dummy (δ^+) is equal to 1 if the share of expected withdrawals is greater than the previous-round share of withdrawals; similarly, when estimating the probability of experimenting with $\boldsymbol{n}$, the dummy (δ^-) is equal to 1 if the share of expected withdrawals is lower than the previous-round share of withdrawals. These dummies highlight whether or not subjects' expectations are aligned with their feedbacks, which may reveal that, for those subjects, the convergence process toward the $(t-1)$-best response can be reversed. Table 4 summarizes main the variables used in our analysis.

4. Results

Our comparative analysis shows that individual withdrawal probability increases with the bank size. In addition to the size of the group, the main drivers of subjects' decisions are the experience via past interactions and the expectations about fellow depositors choices. Furthermore, a non-negligible share of choices are consistent with experimentation towards the non-withdrawal decision, which is more frequent in smaller groups and in earlier rounds.

4.1. Determinants of withdrawals

Table 5 reports the withdrawal rates by bank size and tests the between-treatment differences in the first round via two-sample proportion tests. Although initial choices are similar across bank sizes, the overall withdrawal rates increase with size, and range from 50% in small banks to 88% in large ones.

Fig. 1 reports the distribution of expected withdrawals (via the vertical bars)[21] and shows that expectations are distributed almost symmetrically in small banks, while they are right-skewed in medium and large banks. The figure also reports actual withdrawal rates associated to each expected withdrawal level (via the diamond-dashed line)[22] and shows that the majority of choices are in line with the best response (via the solid line),[23] especially when these are above the threshold q_w. The figure also shows that a sizeable amount of decisions is inconsistent with best response, especially in small banks relative to medium and large ones, and when the expected withdrawal rate is close to q_w. For example, in small banks almost 20% of subjects, whose expected withdrawal rate is lower than q_w, actually choose to withdraw. In medium and large banks this share is quite high (40%), but one should consider that this figure pertains to very few observations. Experimentation may offer a rationale for such behaviors in the neighborhood of the threshold q_w in different groups, which we will analyze via the econometric analysis.

Fig. 2 shows the expected withdrawal rate in the current round ($\texttt{Expectation}_t$), as well as the actual, hence observed via feedback, withdrawal rate in the previous round ($\texttt{Feedback}_{t-1}$), by round and bank size.[24] For all sizes, both rates start at around 40% and exhibit a strong correlation (0.778, p-value < 0.001). However, the dynamic pattern differs across bank size: it is rather flat in small banks, whereas in medium and large banks it steadily increases, so that, in larger banks, it takes less than ten rounds for expected and observed withdrawal rates to be larger than 80%; in these banks, both rates are eventually equal to 100%.

[21] Bars measure the share of k expected withdrawals, $k = 0, \ldots, N-1$ and $N = \{5, 7, 10\}$.

[22] Diamonds measure the share of withdrawals among subjects expecting k withdrawals, $k = 0, \ldots, N-1$ and $N = \{5, 7, 10\}$.

[23] The best response withdrawal rate is equal to 0 (1) if the expected withdrawal rate is (smaller) larger than q_w.

[24] Fig. A.2 in Appendix A shows the dynamics of withdrawal rates in individual banks by session.

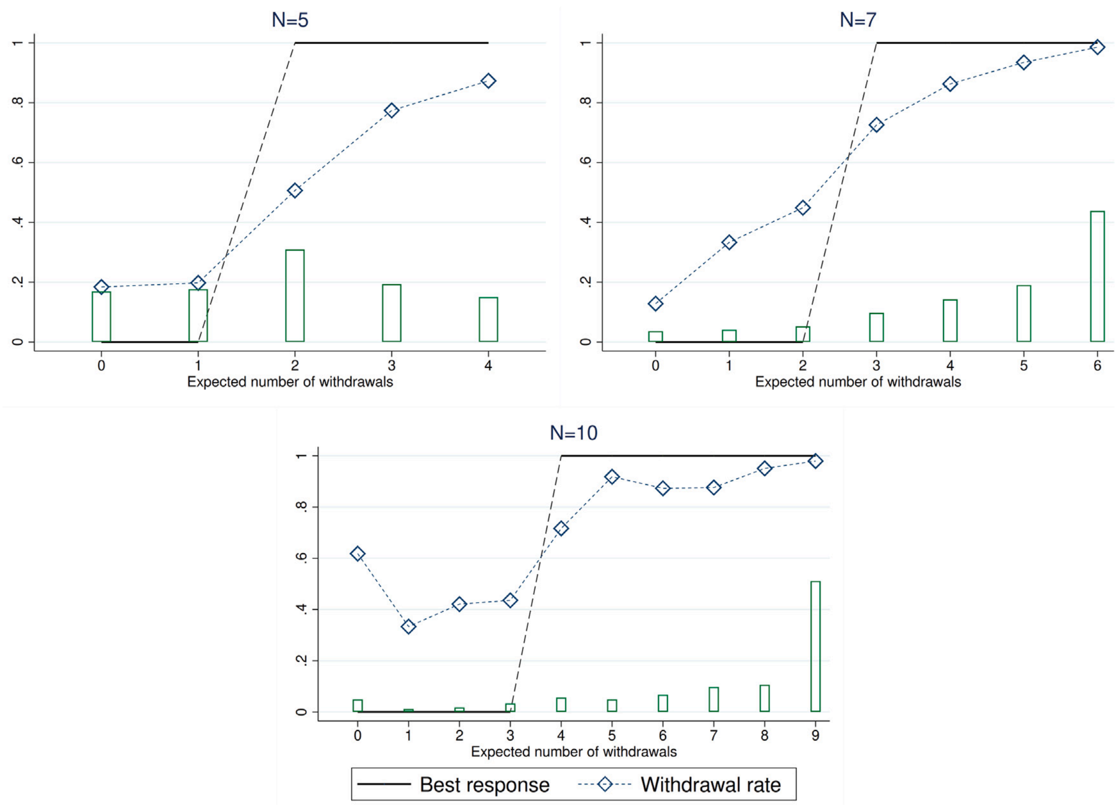

Fig. 1. Expected and actual withdrawal rates.

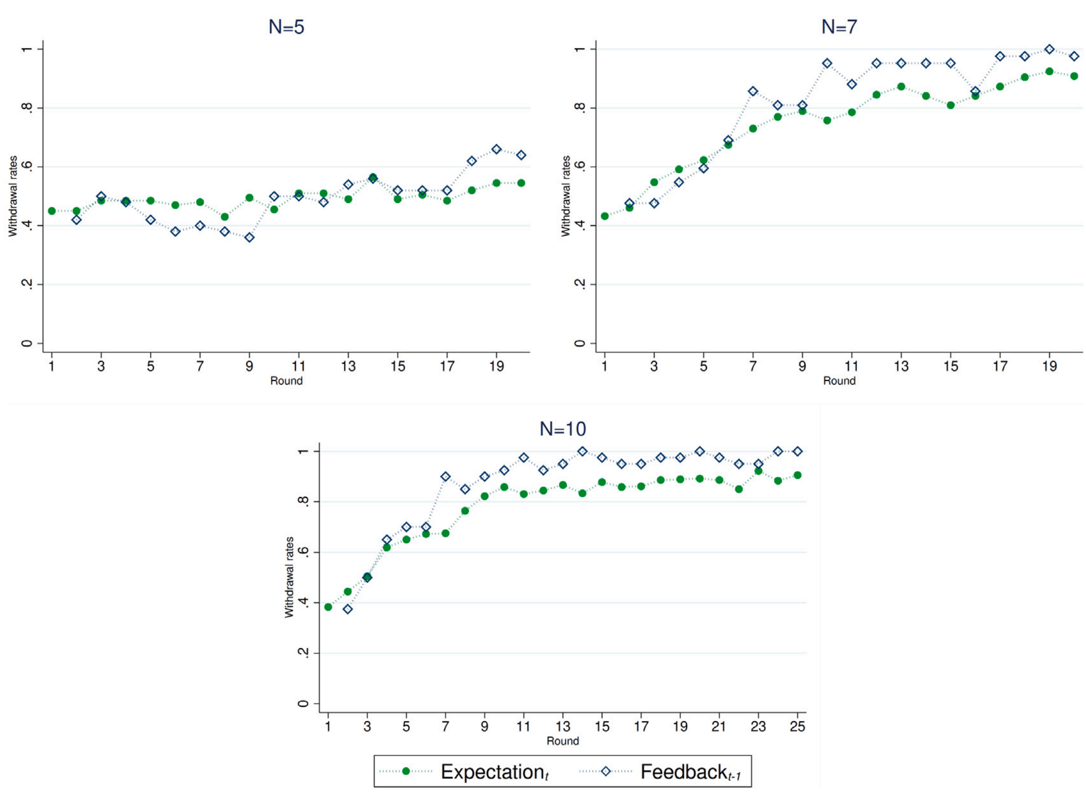

Fig. 2. Expected and observed withdrawal rates.

Table 5
Withdrawal rates by bank size, with proportion tests for first round.

Variable	N = 5		N = 7		N = 10		N = 5 & N = 7	N = 5 & N = 10	N = 7 & N = 10
							Pairwise t-test		
	Obs.	Mean/(SE)	Obs.	Mean/(SE)	Obs.	Mean/(SE)	Mean diff./(p)	Mean diff./(p)	Mean diff./(p)
Withdraw	1000	0.503 (0.016)	840	0.833 (0.013)	1000	0.879 (0.010)			
Withdraw_1	50	0.420 (0.071)	42	0.476 (0.078)	40	0.375 (0.078)	−0.056 (0.589)	0.045 (0.665)	0.101 (0.354)

Mean and standard deviation of withdrawal decision by treatment, for the full sample (Withdraw) and for the first round (Withdraw_1). The latter is tested for between-treatment differences using two-sample proportion tests run on independent observations (significance: *=.1, **=.05, ***=.01).

Since the dynamics of expected and observed withdrawal rates are highly correlated, regardless of bank size and despite the stranger protocol rematching, we use separate models to check how much of the participants' decisions can be explained by their experience due to the repeated interactions, and how much by strategic deliberation based on their expectations. To this end, we estimate the dynamic correlated random-effects probit model in (3)–(4) with the withdrawal decision as dependent variable. Table 6 reports the average marginal effects (AMEs) for four models, which differ in how past interactions and expectations are accounted for, i.e., in the proxies which control for the cross-sectional correlation. All models include: previous- and first-round individual choice (Withdraw_{t-1} and Withdraw_1, respectively); size dummies (N7 and N10, with N5 as reference category); condition (Table D) and grouped round dummies; and demographic controls as in Table 4. Furthermore, Model 1 includes the expected withdrawal rate in the current round (Expectation_t), while Model 2 includes the actual, hence observed, withdrawal rate by fellow depositors in the previous round (Feedback_{t-1}, which in this analysis is centered at the bank size level to avoid collinearity).[25] Model 3 substitutes Feedback_{t-1} with a dummy equal to 1 if Feedback_{t-1} is greater than, or equal to, q_w ($\text{Feedback}_{t-1} \geq q_w$). This specification controls for whether or not to withdraw would have been the optimal choice in the previous round, given the actual decisions of fellow depositors. Model 4 extends Model 2 to disentangle the effects of recent and cumulative past experience, since it includes both Feedback_{t-1} and $\text{History}_{1,t-2}$, which is the share of rounds up to $t-2$ in which $\text{Feedback}_{t-1} \geq q_w$ is equal to 1.[26]

Regarding the bank size, all models confirm, via the significance of the size dummies, that the withdrawal probability increases with the size of the group, as already suggested by the descriptive analysis.[27]

Moreover, Model 1 reveals that the main determinant of individual decision to withdraw is the expected withdrawal rate. In particular, it estimates that an increase by 10% in Expectation_t results in an average increase by 3.6% in the withdrawal probability. According to Model 2, the observed withdrawal rate has a similar effect: an increase by 10% in Feedback_{t-1} results in an average increase by 2.2% in the withdrawal probability. Model 3 shows that the withdrawal probability reacts positively to withdrawing being the best choice in the previous round. Finally, Model 4 shows that both Feedback_{t-1} and $\text{History}_{1,t-2}$ significantly affect the withdrawal probability. Overall, these specifications highlight the crucial role of the beliefs about fellow depositors' behavior. Furthermore, they provide evidence of ex-post rational behavior and show that it is not only the observation of the previous-period outcomes that affects participants' choices, but more broadly the past experience of (relatively high) rates of withdrawals.

Table 6
Determinants of withdrawal decision.

	Model 1	Model 2	Model 3	Model 4
Expectation_t	0.355*** (0.022)			
Feedback_{t-1} (centered)		0.222*** (0.031)		0.193*** (0.032)
$\text{Feedback}_{t-1} \geq q_w$			0.096*** (0.018)	
$\text{History}_{1,t-2}$				0.119*** (0.038)
Withdraw_{t-1}	0.049*** (0.016)	0.105*** (0.020)	0.125*** (0.021)	0.079*** (0.020)
Withdraw_1	0.082*** (0.023)	0.112*** (0.028)	0.111*** (0.028)	0.110*** (0.028)
Bank Size (Baseline N5)				
N7	0.099*** (0.030)	0.217*** (0.028)	0.169*** (0.046)	0.189** (0.041)
N10	0.123*** (0.032)	0.267*** (0.030)	0.206*** (0.052)	0.229** (0.048)
Table D	−0.039* (0.022)	−0.051 (0.036)	−0.059* (0.031)	−0.062* (0.037)
Demographics	✓	✓	✓	✓
Grouped round dummies	✓	✓	✓	✓
Observations	2708	2708	2708	2576
AIC	1630.41	1908.13	1933.39	1735.86
ρ	0.28	0.33	0.31	0.35
LRtest ($H_0 : \rho = 0$)	0.000	0.000	0.000	0.000

Standard errors in parentheses *p < .1; **p < .05; ***p < .01. The dependent variable in all models is the withdrawal decision. The table reports average marginal effects from dynamic correlated random-effect probit models.

Additionally, our estimates confirm the state dependence in the withdrawal decisions: withdrawing in the initial and in the previous rounds unambiguously increases the withdrawal probability. We also find evidence of spurious state dependence, as suggested by the statistically significant contribution of the random effects to the total variance (ρ, see the bottom panel in Table 6). Moreover, consistently with the evidence in Table A.1 and Fig. A.1, the AME of the condition dummy is negative, although weakly significant. Finally, none of the demographic controls, included to account for observable heterogeneity, is consistently significant across the specifications. Overall, all model specifications confirm Hypothesis 1, hence we can state:

Result 1. *The probability to withdraw significantly increases with the size of the bank.*

The Akaike Information Criteria reported in Table 6 (AIC) clearly indicate that Model 1 is the best performing one, we hence run it separately for each group size to complement the previous pooled analysis

[25] The variable assumes numerical values which are typical of each bank size, as shown in Table A.2 in Appendix A. Therefore, its effect would have been confounded by the size dummies. Centering the variable at the bank size level, which we perform using the Stata command center (Jann, 2017), allows to disentangle the true effect of the size from that of the feedback.

[26] For the sake of completeness, in Table A.3 of Appendix A we presents a specification including Expectation_t and Feedback_{t-1}, and this confirms that the fit does not improve significantly and there is multicollinearity.

[27] As already noted in the descriptive analyses (see Fig. 2), 7- and 10-depositors banks exhibit a similar behavior, which is confirmed by the fact that the coefficients associated to N7 and N10 do not differ significantly (Wald test, p-value between 0.184, Model 2, and 0.395, Model 1). Nevertheless, the results clearly highlight a statistically significant difference in withdrawals between small banks (N5) and large ones (N7 and N10).

Table 7
Determinants of withdrawal decision by bank size.

	$N = 5$	$N = 7$	$N = 10$
Expectation$_t$	0.514***	0.322***	0.153***
	(0.051)	(0.030)	(0.030)
Withdraw$_{t-1}$	0.021	0.029	0.056**
	(0.028)	(0.026)	(0.028)
Withdraw$_1$	0.227***	0.056**	−0.010
	(0.074)	(0.025)	(0.024)
Demographics	✓	✓	✓
Grouped round dummies	✓	✓	✓
Observations	950	798	960
Log-likelihood	−388.18	−168.09	−205.80
ρ	0.41	0.07	0.09
LRtest ($H_0 : \rho = 0$)	0.000	0.279	0.099

Standard errors in parentheses *p < .1; **p < .05; ***p < .01. The dependent variable is the withdrawal decision. The table reports average marginal effects from dynamic correlated random-effect probit models, estimated separately for each bank size. The specification is that of Model 1 in Table 6 and includes the expected withdrawal rate, the withdrawal decision in the last round, the initial condition, group of round dummies, and demographic controls, as defined in Table 4.

(see Table 7).[28] These analyses highlight that the initial-condition effect for small and medium banks is more relevant than one's own withdrawal choice in the previous round, whereas the opposite is true for large banks, in which the initial-condition effect vanishes completely and the previous-period choice gains significance. Interestingly, the coefficient of Expectation$_t$ decreases monotonically with bank size, which is consistent with the fact that, in large banks, both expectations and withdrawals quickly reach a very high level and remain persistently high until the end of the experiment (see Fig. 2).[29]

4.2. Determinants of experimentation

In Fig. 3 we report frequencies of experimentation and $(t-1)$-best response by round and bank size. Specifically, the dark-blue area represents ex-post rational choices, i.e., the $(t-1)$-best responses, the light-blue area experimentation with not withdrawing, $e_t(n)$, and the mid-blue area experimentation with withdrawing, $e_t(w)$.

Although ex-post rational behavior is prevalent in all bank sizes, experimentation is not negligible, and it is more frequent in small than in medium and large banks. In initial rounds it counts up to about half of all choices in every group. Subsequently, it remains fairly stable in small banks, while in medium and large ones it declines across rounds, especially in large banks in which subjects' choices rapidly converge to withdrawing. Since in these banks the withdrawal rates are relatively high already in the first rounds and rapidly increase afterwards (see Fig. 2), it seldom happens that $\boldsymbol{n}$ would have been the best choice in round $t-1$. For this reason, in the larger banks experimentation with $\boldsymbol{w}$ rarely occurs and, if at all, only in the initial rounds.

In 231 out of 459 cases in which participants had chosen to experiment with $\boldsymbol{n}$, this choice was also best response given participants' beliefs. Moreover, in 146 out of 164 cases in which participants experimented with $\boldsymbol{w}$, this choice was also best response given their beliefs. While we cannot exclude that some experimentation can be attributed

Table 8
Determinants of experimentation probability.

	$e_{it}(n)$	$e_{it}(w)$
Distance$_{t-1}$ (centered)	−0.339***	0.343*
	(0.043)	(0.180)
δ^-	0.162***	
	(0.015)	
δ^+		0.342***
		(0.047)
Bank Size (Baseline: N5)		
N7	−0.228***	0.176*
	(0.024)	(0.094)
N10	−0.270***	0.296***
	(0.026)	(0.086)
Grouped rounds (Baseline: 1–5)		
6–10	−0.079***	−0.008
	(0.022)	(0.050)
11–15	−0.116***	0.026
	(0.024)	(0.055)
16–20	−0.171***	0.050
	(0.024)	(0.062)
21–25	−0.131***	–
	(0.041)	
Table D	0.038	−0.018
	(0.031)	(0.095)
Demographics	✓	✓
Observations	2321	387
Log-likelihood	1397.82	365.93
ρ	0.32	0.52
LRtest ($H_0 : \rho = 0$)	0.000	0.000

Standard errors in parentheses *p < .1; **p < .05; ***p < .01. Dependent variables are experimentation choices, either with n (first column) or w (second column). The table reports average marginal effects from correlated random-effect probit models. Both models include the distance between Feedback$_{t-1}$ and q_w, a dummy controlling for the sign of the difference between Expectation$_{t-1}$ and Feedback$_{t-1}$, treatment and grouped round dummies, and demographic controls, defined in Table 4. Group of rounds 21 to 25 is omitted in the estimates for $e_{it}(w)$ due to collinearity.

to subjects' mistakes, we take these figures as a first indication that experimentation is consistent with strategic considerations for a sizeable share of participants.

We investigate the determinants of individual experimentation by means of the random-effect probit model in Eqs. (5)–(6) and report its estimates in Table 8.[30] In these models, consistently with our econometric strategy, we use Distance$_{t-1}$, δ^+, and δ^- to control for the cross-sectional correlation, since they are based on Feedback$_{t-1}$. Since the sample size for $e_t(w)$ is admittedly small, we report the results for completeness but, henceforth, focus our comments on the probability to experiment with $\boldsymbol{n}$.

Results confirms that experimentation is more likely in small banks, since its probability diminishes by 22.8% in medium banks and by 27% in large banks, relative to small ones. Furthermore, results in Table 8 reveal that participants are less likely to experiment when the share of withdrawals observed in $t-1$ is farther from q_w (Distance$_{t-1}$),[31] and they are more likely to experiment when they expect less withdrawals than in the previous round (δ^-). The grouped round dummies statistically confirm the evidence in Fig. 3, as the likelihood of experimentation is lower in late than in early (1 to 5) rounds. Lastly, the effect of the condition dummy (Table D) is not statistically significant. Furthermore, Fig. 4 shows that, across all grouped rounds, the experimentation probability is significantly higher when participants expect less withdrawals in the current round ($\delta^- = 1$), relative to when they expect more withdrawals ($\delta^- = 0$), for every bank size.[32]

[28] We exclude the condition dummy Table D from the covariates: since there is one experimental session per combination of bank size and condition, this dummy would possibly capture a session effect, instead of an effect of the condition. For the sake of completeness, we report the results of the size-split model with both Expectation$_t$ and Feedback$_{t-1}$ in Table A.4 in Appendix A.

[29] Fig. A.3 in Appendix A shows how the AMEs of Expectation$_t$ evolve across rounds for each bank size. As the bank size increases a clear time-expectation trade-off emerges: the effect is strong only for low levels of expected withdrawals, which however are observed only for few individuals and in very early rounds.

[30] Observe that, since Distance$_{t-1}$ is a transformation of Feedback$_{t-1}$, we use its centered value at the size level.

[31] This is similar to the finding of Arifovic et al. (2013). In their experiment, however, only some of the participants can identify the $(t-1)$-best response, while in our experiment every participant can do it.

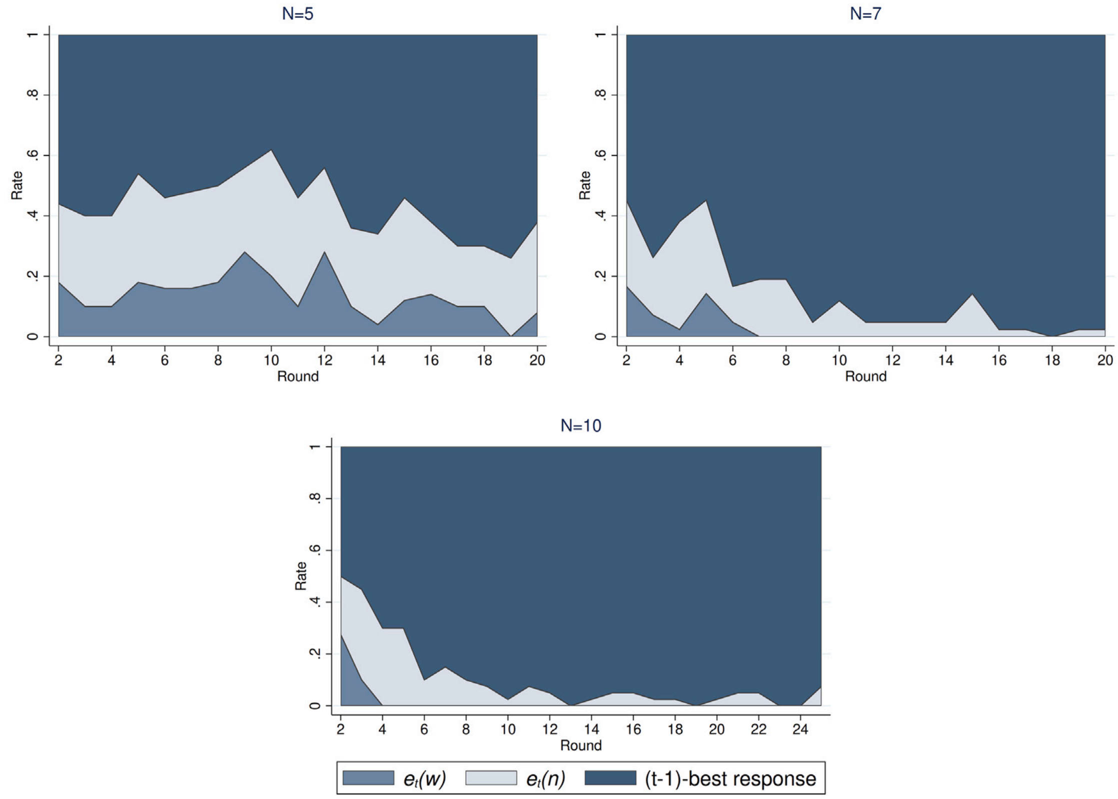

Fig. 3. Experimentation and $(t-1)$-best response dynamics.

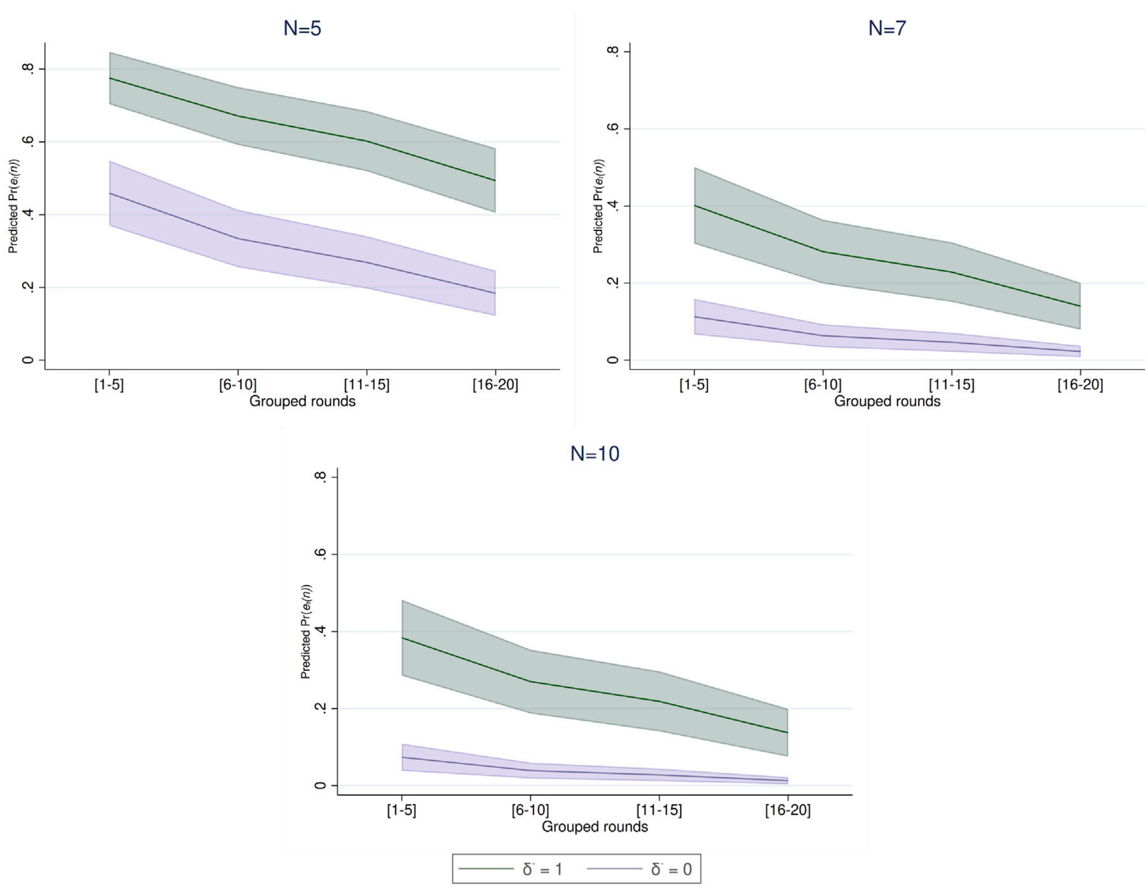

Fig. 4. Dynamics of experimentation probability by δ^-. *Note:* The figures were produced using the Stata command `coefplot`, Jann (2014). The shaded areas correspond to 5% confidence intervals.

Overall, these findings confirm Hypothesis 2. Therefore, we can state:

Result 2. *The probability to experiment with* $\boldsymbol{n}$ *significantly decreases with the size of the bank and it is higher in earlier rounds. Moreover, it is significantly lower when the observed withdrawal rate is farther from* q_w *and greater than the expected withdrawal rate.*

The relationship between experimentation and group size is consistent with the idea that participants acknowledge that their choices can be strategically more relevant in a smaller group, as also noticed by Arifovic et al. (2023). Our results add to this perspective the fact that experimentation with not withdrawing can be interpreted as an attempt of participants to orient their group towards the efficient equilibrium. Such an interpretation is corroborated by the fact that $\texttt{Distance}_{t-1}$ and δ^- are related to the expected individual cost in case the attempt fails. This cost is equal to the difference between the payoff from withdrawing and not withdrawing, given the choices of fellow depositors, and it is increasing with $\texttt{Distance}_{t-1}$ and lower if $\delta^- = 1$. Therefore, our results suggest that experimentation with not withdrawing is more likely when its expected cost is lower.

Since experimentation is not randomly distributed with respect to $\texttt{Distance}_{t-1}$ and δ^-, and these strategic determinants are strongly significant even after learning is accounted for via the grouped round dummies, our findings support the idea that the decision to experiment is grounded on strategic considerations, which prevail over noisy behavior.

5. Conclusion

We examine the role of group size in a bank-run experiment à la (Diamond and Dybvig, 1983), taking a comparative perspective as in Arifovic et al. (2023). Our data show that in small banks the withdrawal rate remains quite stable over repetitions and leads to outcome indeterminacy, whereas in medium and large banks it increases steadily and leads to universal withdrawing in the final rounds. These results are in line with Arifovic et al. (2013) for small banks, but not for larger ones. Indeed, we never observe instances of the efficient outcome in larger groups, an evidence that one can partly attribute to the different protocols: we rely on a stranger protocol while they implement a partner one. Indeed, the re-matching process implies lack of a common history of play among depositors and more uncertain beliefs, which have been shown to reduce the possibility to sustain efficient coordination compared to a fixed-partner protocol (see, e.g. Devetag and Ortmann (2007), and Clark and Sefton (2001)).

Furthermore, we find that bank size is a determinant of the participants' choices and that the individual withdrawal probability is on average 12% higher in large than in small banks. The fact that the withdrawal probability increases with the bank size confirms the results of Arifovic et al. (2023) even in not-so-large groups.

Experience also affects participants' decisions in line with the evidence of Garratt and Keister (2009) and Kiss et al. (2014). In particular, the withdrawal probability increases when participants observe either high withdrawal rates in the previous period or a history of persistently high rates. Along the same lines of the findings in Arifovic et al. (2023), for a fixed coordination tightness, as the dimension of group increases the strategic aspects of the interaction, captured by the subject's expectation about fellow depositors' behaviors, lose their significance in favor of experience. According to the evidence of our experiment, when deciding whether or not to withdraw her deposit from a bank, a participant takes into consideration her experience in other interactions. This points at a potential source of fragility of banks which stems from the fact that the larger the pool of depositors the more likely someone in the pool also holds accounts in other banks. The technological developments of the digital banking services, by reducing the costs of managing several accounts, may increase the pool of depositors and the instances of multiple-banking, and thus make the financial intermediaries more fragile.

We also find that participants engage in experimentation, especially in small banks and in earlier rounds. Our analyses contribute to disentangle the strategic determinants of the decision to experiment. Indeed, the relationship between experimentation and bank size is consistent with participants' perception of being strategically more relevant in smaller groups, as also noticed by Arifovic et al. (2023), and with their weighing the effectiveness of this choice against its expected cost. In addition, the finding that experimentation is non-randomly distributed with respect to variables linked to feedback and to beliefs, even after learning across rounds is accounted for, supports the idea that this decision is based on strategic considerations.

CRediT authorship contribution statement

Federico Belotti: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Eloisa Campioni:** Writing – review & editing, Writing – original draft, Visualization, Validation, Resources, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Vittorio Larocca:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Francesca Marazzi:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Luca Panaccione:** Writing – review & editing, Writing – original draft, Visualization, Validation, Resources, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Andrea Piano Mortari:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Additional tables and figures

See Tables A.1–A.4 and Figs. A.1–A.4.

Appendix B. Supplementary material

Payoff tables, condition Table I

See Tables B.1–B.3.

Payoff tables, condition Table D

See Tables B.4–B.6.

[32] Fig. A.4 in Appendix A reports the dynamics of $\texttt{Distance}_{t-1}$ and of the experimentation probability predicted by our model, and confirms that as $\texttt{Distance}_{t-1}$ increases across rounds, the experimentation probability fades out, again at a faster rate in medium and large banks.

Table A.1
Withdrawal rates by condition, with proportion tests for first round.

Variable	Table I		Table D		Table I & Table D Pairwise t-test
	Obs.	Mean/(SE)	Obs.	Mean/(SE)	Mean diff./(p)
Withdraw	1420	0.773 (0.011)	1420	0.694 (0.012)	
$Withdraw_1$	66	0.455 (0.062)	66	0.394 (0.061)	0.061 (0.481)

Mean and standard deviation of withdrawal decision by condition, for the full sample (Withdraw) and for the first round ($Withdraw_1$). The latter is tested for differences between conditions using a two-sample proportion test run on independent observations (significance: *=.1, **=.05, ***=.01).

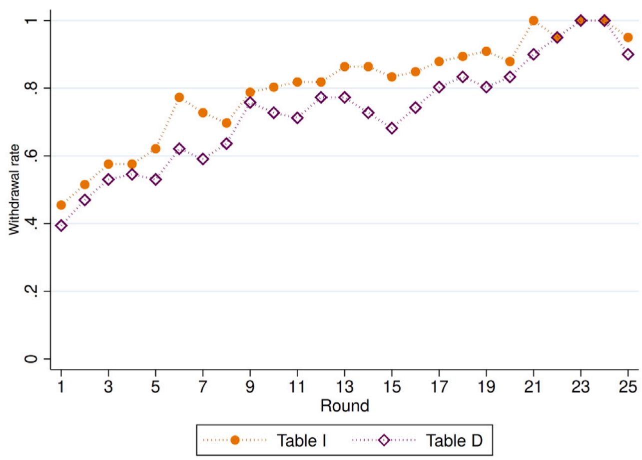

Fig. A.1. Withdrawal rates by round and condition.

Table A.3
Determinants of withdrawal decision with $\texttt{Expectation}_t$ and $\texttt{Feedback}_{t-1}$ as explanatory variables.

	Model 1A
$Expectation_t$	0.335*** (0.022)
$Feedback_{t-1}$ (centered)	0.118*** (0.028)
$Withdraw_{t-1}$	0.045*** (0.016)
$Withdraw_1$	0.084*** (0.023)
Bank Size (Baseline: N5)	
N7	0.098*** (0.030)
N10	0.125*** (0.032)
Table D	−0.020 (0.029)
Demographics	✓
Grouped round dummies	✓
Observations	2708
AIC	1615.39
ρ	0.28
LRtest ($H_0 : \rho = 0$)	0.000

Standard errors in parentheses.
*p < .1; **p < .05; ***p < .01.

Table A.2
Distribution of $\texttt{Feedback}_{t-1}$ by bank size.

	N = 5 %	N = 7 %	N = 10 %	Total %
0	7.8	0.2	0.0	2.8
.11			0.0	0.0
.17		1.9		0.6
.22			0.3	0.1
.25	24.4			8.6
.33		3.0	2.3	1.7
.44			5.4	1.9
.50	35.4	9.0		15.1
.56			4.2	1.5
.67		13.3	5.0	5.7
.75	23.6			8.3
.78			7.4	2.6
.83		23.3		6.9
.89			19.4	6.8
1	8.8	49.2	56.0	37.4
Total	100.0	100.0	100.0	100.0

Table A.4
Determinants of withdrawal decision by bank size with $\texttt{Feedback}_{t-1}$ (Panel A) and both $\texttt{Feedback}_{t-1}$ and $\texttt{Expectation}_t$ (Panel B) as explanatory variables.

Panel A: $\texttt{Feedback}_{t-1}$			
	$N = 5$	$N = 7$	$N = 10$
$Feedback_{t-1}$	0.267*** (0.053)	0.060 (0.067)	0.072 (0.061)
$Withdraw_{t-1}$	0.085*** (0.030)	0.076*** (0.028)	0.064*** (0.022)
$Withdraw_0$	0.266*** (0.064)	0.087** (0.036)	−0.024 (0.028)
Panel B: $\texttt{Feedback}_{t-1}$ and $\texttt{Expectation}_t$			
$Expectation_t$	0.486*** (0.050)	0.320*** (0.029)	0.152*** (0.030)
$Feedback_{t-1}$	0.158*** (0.048)	−0.013 (0.057)	0.014 (0.060)
$Withdraw_{t-1}$	0.025 (0.028)	0.029 (0.026)	0.056** (0.029)
$Withdraw_1$	0.225*** (0.071)	0.045* (0.024)	−0.009 (0.024)
Demographics	✓	✓	✓
Grouped round dummies	✓	✓	✓
Observations	950	798	960

Standard errors in parentheses. *p < .1; **p < .05; ***p < .01.

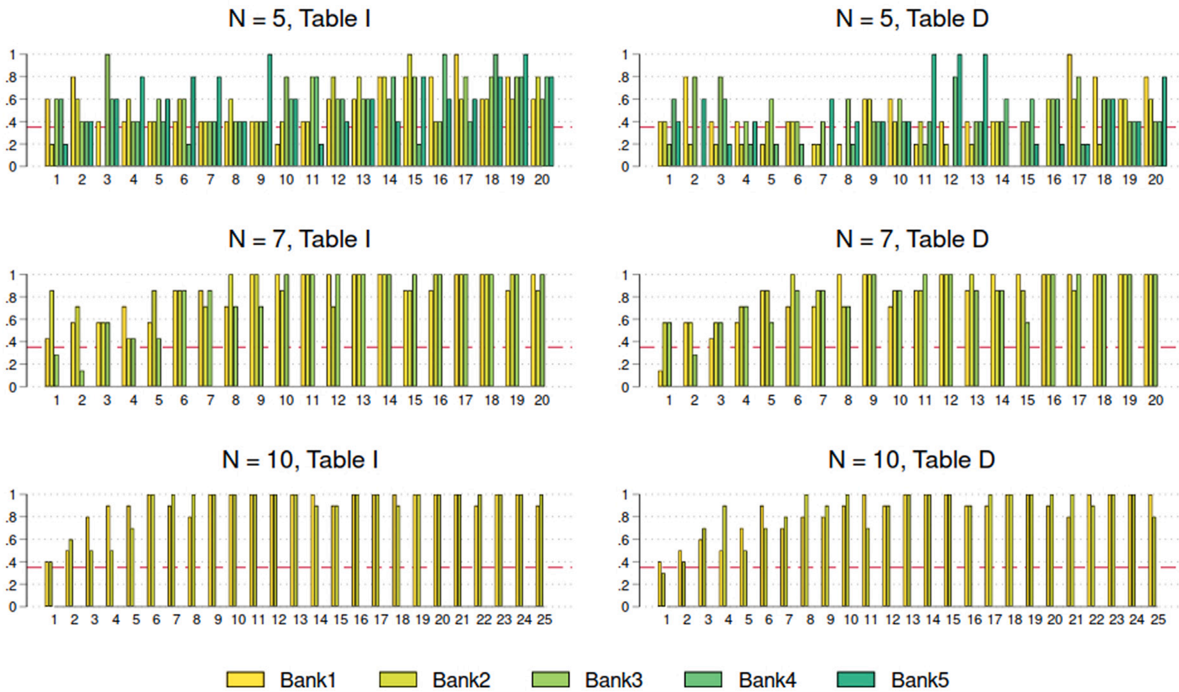

Fig. A.2. Dynamics of withdrawal rates in individual banks by session. *Note:* the red dashed line indicates the threshold q_w.

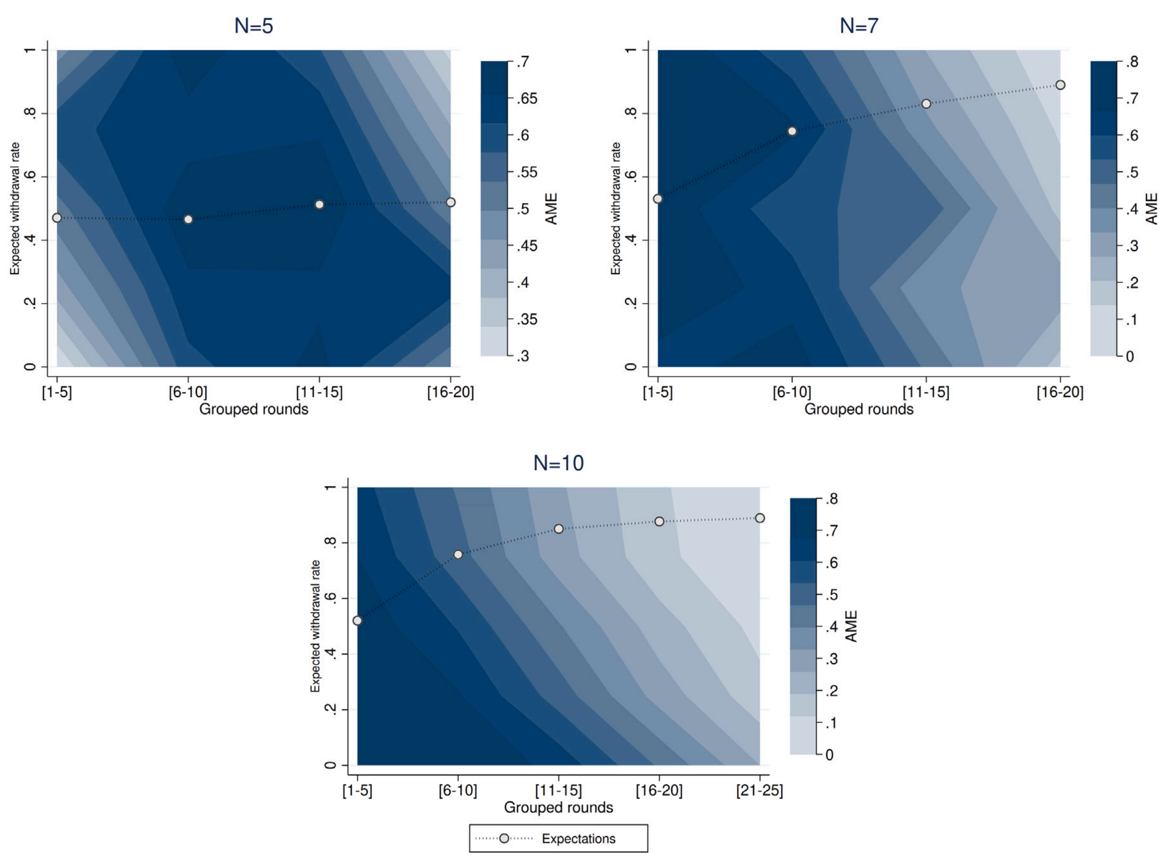

Fig. A.3. Average marginal effect of `Expectations` across rounds. *Note:* plots were produced using the Stata command `twoway contour` and show how the average marginal effects (AMEs) of `Expectation` on `Withdraw` changes across (group of) rounds and by bank size. The intensity of the color represents the magnitude of the AME, which are estimated for every combination of group of rounds and expectation value. The white dots represent the average expectations for a given block of rounds.

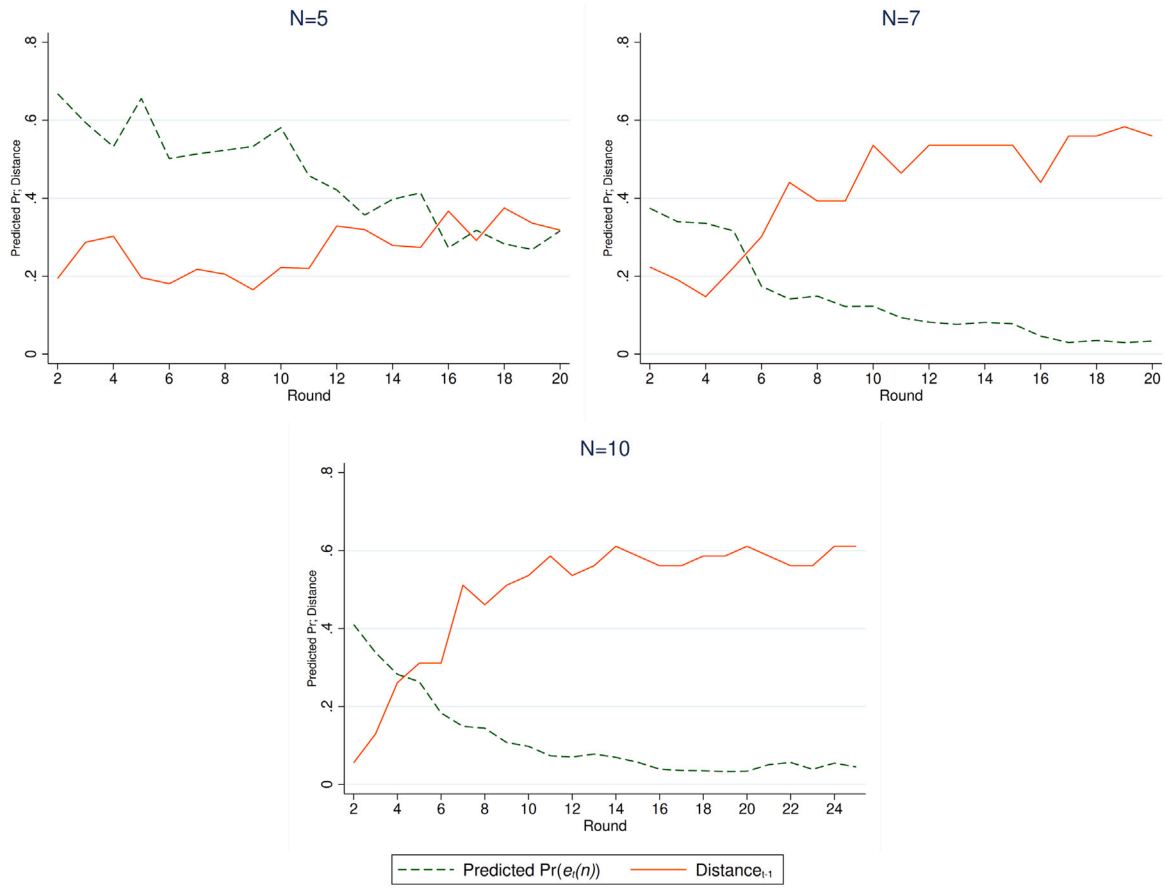

Fig. A.4. Predicted experimentation probability and distance from threshold over rounds. *Note:* The figure reports the dynamics of Distance_{t-1} (continuous, red line) and of the experimentation probability (dashed, green line) predicted by our model across all rounds and by bank size.

Table B.1
Payoff matrix for banks with $N = 5$, condition *Table I*.

	Payoff if you withdraw ∘	Payoff if you do not withdraw •
∘∘∘∘	98	7
•∘∘∘	122	90
••∘∘	122	117
•••∘	122	132
••••	122	150

Table B.2
Payoff matrix for banks with $N = 7$, condition *Table I*.

	Payoff if you withdraw ∘	Payoff if you do not withdraw •
∘∘∘∘∘∘	98	7
•∘∘∘∘∘	117	57
••∘∘∘∘	122	96
•••∘∘∘	122	115
••••∘∘	122	127
•••••∘	122	134
••••••	122	150

Table B.3
Payoff matrix for banks with $N = 10$, condition *Table I*.

	Payoff if you withdraw ∘	Payoff if you do not withdraw •
∘∘∘∘∘∘∘∘∘	99	0
•∘∘∘∘∘∘∘∘	117	7
••∘∘∘∘∘∘∘	122	63
•••∘∘∘∘∘∘	122	90
••••∘∘∘∘∘	122	107
•••••∘∘∘∘	122	118
••••••∘∘∘	122	126
•••••••∘∘	122	132
••••••••∘	122	136
•••••••••	122	150

Table B.4
Payoff matrix for banks with $N = 5$, condition *Table D*.

	Payoff if you withdraw ∘	Payoff if you do not withdraw •
••••	122	150
•••∘	122	132
••∘∘	122	118
•∘∘∘	122	90
∘∘∘∘	98	7

Table B.5
Payoff matrix for banks with $N = 7$, condition *Table D*.

	Payoff if you withdraw ○	Payoff if you do not withdraw •
••••••	122	150
•••••○	122	134
••••○○	122	127
•••○○○	122	115
••○○○○	122	96
•○○○○○	117	57
○○○○○○	98	7

Table B.6
Payoff matrix for banks with $N = 10$, condition *Table D*.

	Payoff if you withdraw ○	Payoff if you do not withdraw •
•••••••••	122	150
••••••••○	122	136
•••••••○○	122	132
••••••○○○	122	126
•••••○○○○	122	118
••••○○○○○	122	107
•••○○○○○○	122	90
••○○○○○○○	122	63
•○○○○○○○○	117	7
○○○○○○○○○	99	0

Instructions[33]

Introduction

Welcome! You are participating in an experiment to collect data for a scientific research.

During the experiment you have to make decisions that will contribute to determine a payoff, that will be paid cash at the end of the experiment.

The experiment is totally anonymous: neither the experimenters nor other participants will be able to associate your decisions to your identity.

During the experiment, your interactions with other participants will be intermediated by a computer. Any form of communication between participants is prohibited. If you violate this rule, you will be excluded from the experiment with no payment.

If you have any doubt about the experiment, raise your hand and an experimenter will come to answer to your question, privately.

The experiment consists of a sequence of several phases. For each of phase, you will receive specific instructions.

All the decisions you will take in each phase will contribute to your final payoff. In some phases your payoff depends only on your own decisions, while in others it depends on your decisions and on the decisions of other participants, as it will be explained later on.

Your payoff in each phase and your final payoff are expressed in an experimental currency called *Zed*. Your final payoff in Zed will be converted into a final payment in Euros, at the exchange rate of 20 Zed = 1 Euro.

Phase 1

In this phase you will be asked to answer to 13 questions. Every question has four possible answers, and your task is to choose the correct answer. For every question there is only one correct answer. You must answer the questions on your own and your payoff depends only on your choices. For each correct answer you will receive 1 point whereas for any wrong answer you will lose 1/2 points.

[33] This represents a translation of the instructions used in the experiment for the groups of five depositors. The actual instructions are in Italian.

The questions will appear sequentially on your screen, and for each question you have 90 s to answer. If you do not provide any answer within the given time, that question will be considered as unanswered and you will not gain nor lose any point. *Please note that once provided, your answer cannot be changed.*

At the end of this phase, the computer screen will summarize: your own answers, the correct answers and the points you gained.

Your payoff in Phase 1

Your payoff in this phase depends on your answers to the questionnaire and on a binary lottery that guarantees a prize of 150 Zed or of 50 Zed.

The total points obtained from the questionnaire will determine the probability of winning the prize of 150 Zed. This probability cannot be lower than 0 nor greater than 1, and it increases with the points obtained. Recall that the probability of gaining the prize of 50 is one minus the probability of gaining the prize of 150.

If all your answers are wrong, your score from the questionnaire is $(-1/2) \times 13 = -6.5$ and the probability of the prize of 150 is equal to 5%: this is the lowest probability with which you can win the high prize (150 Zed). In this case, the probability of the prize of 50 Zed is equal to 95%.

On the other hand, if all your answers are correct, then your score is $1 \times 13 = 13$ and the probability the prize of 150 Zed is 95%: this is the highest probability with which you can win the high prize. In this case, the probability of the prize of 50 Zed is equal to 5%.

For any other score, you will win the prize of 150 Zeds with a probability between 5% and 95%.

The lottery draw over the two prizes will be performed at the end of the experiment and the prize will be part of your final payoff.

For your convenience, we are providing you with a blank table that you can use to take note of the results of the questionnaire.

Phase 2

Let us now move to Phase 2.

An experimenter will read aloud the instructions of this phase. If you have any question, please raise your hand and an experimenter will come to answer your question, privately.

Recall that communication between participants is prohibited. If you violate this rule, you will be excluded from the experiment with no payment.

Your task in Phase 2

In this phase, you and other 4 participants will be randomly and anonymously selected to constitute an experimental bank.

Every member of the bank owns 100 Zed deposited in the experimental bank. Hence, a bank is composed of 5 depositors, whose identity is unknown to each other.

As a depositor, you have two options: you can either withdraw your 100 Zed and close your deposit account; or you can leave your money deposited in the bank.

How much you receive in either case depends jointly on how much the bank promises to repay and on the decisions of the depositors at your bank, who face your identical task.

The bank promises to repay 150 Zed to every depositor who decides not to withdraw his money and 122 Zed to every depositor who decides to withdraw. However, the bank may not be able to fulfill her promises if too many depositors decide to withdraw. The table below lists the payoffs you obtain depending on your choice and on the choices of all other depositors in your bank.

Payoff Table		
	Payoff if you withdraw ∘	Payoff if you do not withdraw •
∘ ∘ ∘ ∘	98	7
• ∘ ∘ ∘	122	90
• • ∘ ∘	122	117
• • • ∘	122	132
• • • •	122	150

The bullets in the first column represent the possible decisions of the depositors at your bank other than you. In particular, the white bullet represents a depositor who decided to withdraw and close his deposit. The black bullet, on the contrary, represents a depositor who decided not to withdraw.

Example 1. Suppose that all depositors other than you withdraw. As the table shows, if you withdraw your payoff is 98 Zed. If you do not withdraw, your payoff is 7 Zed (see the first row of the table).

Example 2. Suppose that 3 depositors other than you decide not to withdraw. As the table shows, if you withdraw your payoff is 122 Zed. If you do not withdraw, your payoff will be 132 Zed (see the fourth of the table).

Why cannot the bank always guarantee the promised repayments? Imagine that once the experimental bank has been constituted, the total deposits of 500 Zed are invested and that it takes time to generate a return.

To repay a depositor who decides to withdraw, the bank has to prematurely liquidate part of the investment. Those who do not withdraw are paid with the resources left after having repaid those who withdraw. Since premature liquidation is costly, if too many depositors decide to withdraw the bank cannot guarantee the promised repayments.

At the time you make your choice, the decision of the other depositors is unknown to you. Since any form of communication is forbidden, you are not allowed to ask to other participants their choice.

Procedure for Phase 2

Phase 2 consists of 20 periods. Each period is independent and completely separate from the others. In every period you will perform the task described in the previous section.

In each period, several experimental banks will be constituted, and each of them is completely separate from the others. Depositors are randomly assigned to an experimental bank. Therefore, you will meet with different depositors in every period. We cannot exclude that you will meet the same depositor more than once. However, the assignment to an experimental bank is completely anonymous, hence it is not possible for you to identify the other depositors.

At the beginning of each period you will have 100 Zed deposited in your experimental bank. As a preparation to your main decision, you will be asked to state your expectations about how many depositors of your bank other than you will withdraw, and about how many will leave their money deposited in the bank. Note that the sum of these two numbers has to be equal to 4 (four).

Then, you will have to decide whether to withdraw or not your deposit. You have 30 s to take your decision. If you have not made any decision within the time limit, the computer will randomly select your decision.

At the end of each period your decision, your payoff and the number of withdrawals at your bank will be privately communicated to you.

Computer instructions

During phase 2, three different screens will appear on your computer: preliminary, decision and report screens.

The preliminary screen gives you information about the experimental bank you have been assigned to.

The decision screen shows the payoff table as described above. It shows the two buttons that you will have to press to take your decision.

Once you press a button, you cannot change your choice.

At the top of the screen, the current period is displayed. At the bottom, there is countdown bar showing the time left to take you decision.

After all depositors in your bank have taken their decisions, a report screen will provide information about: your decision, your payoff and the number of depositors who decided to withdraw in the current period. You have 10 s to read those information before the new period starts.

Your payoff in Phase 2

Your payoff for Phase 2 will be determined by random selection of one period out of the 20 ones. The draw will be performed at the end of the experiment and you will be assigned the payoff corresponding to the selected period.

Phase 3

An experimenter will read aloud the instructions of this phase. If you have any question, please raise your hand and an experimenter will come to answer your question, privately.

Recall that communication between participants is prohibited. If you violate this rule, you will be excluded from the experiment with no payment.

In this phase 10 (ten) pairs of binary lotteries will be displayed on your screen. For each lottery, you will find on the screen the value and the probability of each prize. Your task is to choose one lottery within each pair. Your choice will determine your payoff for this phase as described below.

Your payoff in Phase 3

At the end of the experiment, one of the ten lottery pairs will be randomly chosen. Right after, the lottery you chose within the selected pair will be played by your computer. The prize extracted will determine your payoff in Zed for this phase.

Concluding the experiment

This phase is devoted to determine your total payoff, that is the sum of the payoffs you gained in each phase of the experiment.

We start with Phase 1. The computer will summarize on your screen: your answers to the questionnaire, the correct answers, your total score, and your probability to win the prize of 150 Zed. The lottery draw will be visualized on your computer and it will determine your payoff for Phase 1.

As for Phase 2, one period out of the 20 will be randomly chosen. The random draw is common to all participants. At this stage, the computer screen will summarize your payoffs for every period of Phase 2.

As for Phase 3, we will select one of the ten pairs of lotteries through a random procedure. Subsequently, the computer will play the lottery you choose within the selected pair. The prize visualized on your computer will determine your payoff for Phase 3.

The sum of all payoffs will determine your final payoff expressed in Zed. This payoff will then be converted into euro according to the predetermined exchange rate of 20 Zed = 1 Euro, and this amount will constitute your final payment for the experiment.

References

Allen, F., Gale, D., 1998. Optimal financial crises. J. Finance 53 (4), 1245–1284.

Allen, F., Gale, D., 2009. Understanding financial crises. OUP Catalogue, Oxford University Press.

Andolfatto, D., Nosal, E., Sultanum, B., 2017. Preventing bank runs. Theor. Econ. 12 (3), 1003–1028.

Arifovic, J., Hommes, C., Kopányi-Peuker, A., Salle, I., 2023. Ten isn't large! Group size and coordination in a large-scale experiment. Am. Econ. J. Microecon. 15 (1), 580–617.

Arifovic, J., Jiang, J.H., Xu, Y., 2013. Experimental evidence of bank runs as pure coordination failures. J. Econom. Dynam. Control 37 (12), 2446–2465.

Barr, M.S., 2023. Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank. Board of Governors of the Federal Reserve System, Available online.

Calomiris, C.W., Mason, J.R., 2003. Consequences of bank distress during the great depression. Amer. Econ. Rev. 93 (3), 937–947.

Clark, K., Sefton, M., 2001. Repetition and signalling: experimental evidence from games with efficient equilibria. Econom. Lett. 70 (3), 357–362.

Dávila, E., Goldstein, I., 2023. Optimal deposit insurance. J. Polit. Econ. 131 (7), 1676–1730.

Devetag, G., Ortmann, A., 2007. When and why? A critical survey on coordination failure in the laboratory. Exp. Econ. 10 (3), 331–344.

Devlin, J., Gerrard, P., 2005. A study of customer choice criteria for multiple bank users. J. Retail. Consum. Serv. 12 (4), 297–306.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance, and liquidity. J. Polit. Econ. 401–419.

Duffy, J., 2016. Macroeconomics: A survey of laboratory experiments. In: Kagel, J., Roth, A. (Eds.), The Handbook of Experimental Economics. Vol. 2, Princeton University Press, Princeton, NJ, pp. 1–90.

Dufwenberg, M., 2015. Banking on experiments? J. Econ. Stud. 42 (6), 943–971.

Ennis, H.M., Keister, T., 2009. Bank runs and institutions: The perils of intervention. Amer. Econ. Rev. 99 (4), 1588–1607.

Evans, K., 2015. Playing the Field. Consumers and Competition in Banking. Technical Report, Social Market Foundation.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.

Garratt, R., Keister, T., 2009. Bank runs as coordination failures: An experimental study. J. Econ. Behav. Organ. 71 (2), 300–317.

Gerrard, P., Cunningham, B.J., 2001. Singapore's undergraduates: how they choose which bank to patronise. Int. J. Bank Mark. 19 (3), 104–114.

Goldstein, I., Pauzner, A., 2005. Demand-deposit contracts and the probability of bank runs. J. Finance 60 (3), 1293–1327.

Green, E.J., Lin, P., 2003. Implementing efficient allocations in a model of financial intermediation. J. Econom. Theory 109 (1), 1–23.

Greiner, B., 2015. Subject pool recruitment procedures: Organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1 (1), 114–125.

Heckman, J., 1981. Heterogeneity and state dependence. In: Studies in Labor Markets. National Bureau of Economic Research, Inc, pp. 91–140.

Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. Amer. Econ. Rev. 92 (5), 1644–1655.

Jann, B., 2014. Plotting regression coefficients and other estimates. Stata J. 14 (4), 708–737.

Jann, B., 2017. Center: Stata module to center (or standardize) variables. Stat. Softw. Compon. (S4444102).

Kiss, H.J., Rodriguez-Lara, I., Rosa-García, A., 2012. On the effects of deposit insurance and observability on bank runs: An experimental study. J. Money Credit Bank. 44 (8), 1651–1665.

Kiss, H.J., Rodriguez-Lara, I., Rosa-García, A., 2014. Do social networks prevent or promote bank runs? J. Econ. Behav. Organ. 101, 87–99.

Kiss, H.J., Rodriguez-Lara, I., Rosa-García, A., 2016. Think twice before running! Bank runs and cognitive abilities. J. Behav. Exp. Econ. 64, 12–19.

Kiss, H.J., Rodriguez-Lara, I., Rosa-García, A., 2018. Panic bank runs. Econom. Lett. 162, 146–149.

Kiss, H.J., Rodriguez-Lara, I., Rosa-García, A., 2022. Who withdraws first? Line formation during bank runs. J. Bank. Financ. 140, 106491.

Madiès, P., 2006. An experimental exploration of self-fulfilling banking panics: their occurrence, persistence, and prevention. J. Bus. 79, 1831–1866.

OECD, 2012a. PISA 2012 financial literacy questions and answers.

OECD, 2012b. PISA 2012 released creative problem solving items.

OECD, 2012c. PISA 2012 released mathematics items.

OECD, 2014. PISA 2014 financial literacy questions and answers.

OECD, 2015. PISA 2015 released field trial cognitive items.

Peck, J., Shell, K., 2003. Equilibrium bank runs. J. Polit. Econ. 111 (1), 103–123.

Peia, O., Vranceanu, R., 2019. Experimental evidence on bank runs with uncertain deposit coverage. J. Bank. Financ. 106, 214–226.

van Rooij, M., Lusardi, A., Alessie, R., 2011. Financial literacy and stock market participation. J. Financ. Econ. 101 (2), 449–472.

Saunders, A., Wilson, B., 1996. Contagious bank runs: evidence from the 1929–1933 period. J. Final. Intermediation 5 (4), 409–423.

Schotter, A., Yorulmazer, T., 2009. On the dynamics and severity of bank runs: An experimental study. J. Final. Intermediation 18 (2), 217–241.

Selten, R., Abbink, K., Cox, R., 2005. Learning direction theory and the winner's curse. Exp. Econ. 8, 5–20.

Temzelides, T., 1997. Evolution, coordination, and banking panics. J. Monetary Econ. 40 (1), 163–183.

Weimann, J., Brosig-Koch, J., 2019. Methods in experimental economics. Springer Texts in Business and Economics, Springer International Publishing, Berlin.

Wooldridge, J.M., 2005. Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. J. Appl. Econometrics 20 (1), 39–54.