Online Appendix for

"Bank Run, Interrupted: Modeling Deposit Withdrawals with Generative AI"

Abstract

This Internet Appendix provides additional material on the data, methodology, and some robustness exercises related to the main text.

A Supplemental Tables

A.1 Survey Instrument

This section describes the exact survey instrument used in the experiment. Respondents first answered demographic items (household income, education), then viewed a common stem describing a bank-run scenario with a viral tweet, followed by one randomized treatment (Baseline, Network Effect, Bank Email, Bank Text, President Message, Fed Message, Cautionary Tale, or Placebo). After selecting a binary withdrawal decision (Yes/No), they provided a free-text explanation. The instrument was fielded in March 2025 to a representative U.S. sample of 1,200 participants with random assignment across the eight arms.

Q1: Household Income

What was your total annual household income before taxes?

- ◦ Less than $25,000
- ◦ $25,000–$49,999
- ◦ $50,000–$74,999
- ◦ $75,000–$119,999
- ◦ $120,000 or more

Education

What is the highest level of education you have completed?

- ◦ No high school (i.e., did not complete high school)
- ◦ High school diploma or GED
- ◦ Some college or associate degree
- ◦ Bachelor's degree (BA/BS)
- ◦ Graduate or professional degree, such as Masters, Ph.D., J.D., or M.D.

Common Stem (Shown in All Conditions)

Survey Body

Imagine the following scenario: You just heard about a run on a local bank (not yours) in your state.

You go online and see the following tweet from a popular user.

Run on the bank!

Get your money out.

First thing on Monday.

US banks are in trouble.

FED emergency meeting.

Deposits may get locked.

Possible withdrawal limits.

Cash is king. Get out now!

You have some uninsured deposits in your bank (i.e., not covered by FDIC).

[Treatment Arm]

Do you plan to **withdraw your money** from your bank given this information?

Treatment Arms (Randomized)

Baseline

[Empty]

Network Effect

You then run into a friend who says they withdrew funds from *your* bank given the news.

Bank Email

You receive an email from your bank addressing the rumors: "We are aware of the recent rumors circulating on social media. Please be assured that our bank is financially strong, and your funds are safe with us. We pride ourselves on upholding the highest standards of financial integrity, consistently validated by rigorous internal and external audits. Your trust is paramount to us, and we are here to address any concerns you may have."

Bank Text

You receive a text from your contact at the bank: "Just a heads up: despite the online buzz, the bank is doing just fine. Always here if you want to chat."

President Message

You read a news article headlined: "President Affirms: US Banks Remain Stable and Secure."

Fed Message

You read a news article headlined: "Fed Chairman Confirms: Our Banking System is Strong and Stable."

Cautionary Tale

Your regional Reserve Bank tweets about a new article: "The Devastating Impact of Bank Runs during the Great Depression."

Placebo

You see a news article headlined: "A mild and dry start to November."

Explanation

Explain the reason for your answer:

Response type: Free-text (unlimited characters); displayed after the binary withdrawal decision.

A.2 Survey Weights

Table A.1: Synthetic (Gemma) vs. Human (Prolific) vs. U.S. Population

Demographic	Category	Synthetic (%)	Prolific (%)	U.S. (%)
Age				
	18–24	20.00	11.92	9.9
	25–34	20.00	17.10	13.3
	35–54	20.00	32.64	26.2
	55–64	20.00	24.70	12.4
	65+	20.00	13.64	16.5
Gender				
	Man	50.00	49.14	49.2
	Woman	50.00	50.86	50.8
Income				
	$25K–$50K	20.00	22.71	23.0
	$50K–$75K	20.00	17.88	18.0
	$75K–$120K	20.00	21.85	20.0
	Over $120K	20.00	20.90	17.0
	Under $25K	20.00	16.67	22.0
Education				
	College	25.00	40.85	29.0
	Graduate	25.00	23.92	32.0
	High School	25.00	34.20	27.0
	No High School	25.00	1.04	12.0

Notes: Synthetic respondents are generated via Gemma; Prolific figures are observed; U.S. targets are ACS 2023 estimates.

A.3 Simulation Parameterization and Calibration

This subsection summarizes the parameters used in the simulations. Table A.2 provides an overview, while the notes below detail how each component is implemented.

Parameter	Meaning	Value or Source
α_t	Treatment log–odds shift	Estimated from LLM–human calibration
β	Activation probability	0.4
q_0, κ, α	Private precision, gain, cost	$q_0 = 0.6$, $\kappa = 0.6$, $\alpha = 0.6$
λ, η, ζ_0	Survival signal parameters	$\lambda = 0.5$, $\eta = 5$, $\zeta_0 = 1$; $\zeta_\tau = \max\{1 - 0.1(\tau - 1), 0\}$
γ (ramp)	Peer coefficient	$\gamma = 1.0$; ramp over first $R = 3$ rounds
$\bar{\theta}, \sigma_\theta$	Threshold distribution	$\bar{\theta} = 0.56$, $\sigma_\theta = 0.18$ (truncated to $(0, 1)$)
Network	Watts–Strogatz topology	$n = 600$, $k = 8$, $p = 0.1$
Stopping	Convergence criterion	2 no-change rounds; max 30 rounds

Table A.2: Simulation parameters and their sources. Only α_t varies across treatments; all other parameters are fixed.

Treatment shifts, denoted by α_t, are estimated using the error-corrected LLM–human calibration process and vary across treatment conditions. All other parameters are held fixed unless otherwise specified. The simulated network follows a Watts–Strogatz small world topology consisting of $n = 600$ agents, with each agent having a mean degree of $k = 8$ and a rewiring probability of $p = 0.1$. These settings reflect the default configuration used in the runner but can be modified if needed.

Agents become active through a geometric arrival process, with a per-round activation probability of $\beta = 0.4$. Upon becoming active, agents receive private signals whose precision depends on the level of effort exerted. Signal precision is modeled as $q_{i\tau} = q_0 + \kappa a_{i\tau}$, with a quadratic cost of effort $\chi(a) = \alpha a^2$. Here, $\kappa = 0.6$ and $\alpha = 0.6$ are fixed, and the baseline signal precision q_0 is calibrated using error-corrected LLM output. Agents choose effort myopically at each time step, optimizing according to $a_{i\tau} = \min\left\{\frac{\kappa|0.5 - b_i(\tau)|}{2\alpha}, 1\right\}$.

In addition to private signals, agents receive a public survival signal defined by $S_L(\tau) = \frac{e^{\lambda\zeta_\tau} - 1}{e^{\lambda\eta} - 1}$, where $\zeta_\tau = \max \zeta_0 - 0.1(\tau - 1), , 0$, and the parameters are set to $\lambda = 0.5$, $\eta = 5$, and $\zeta_0 = 1$. Peer effects are incorporated through a logit update of the agent's belief: $\operatorname{logit} b_i(\tau + 1) = \operatorname{logit} b_i(\tau) +$

$\gamma_\tau s_i(\tau)$, where $s_i(\tau)$ is the observed share of neighbors who have withdrawn. Peer sensitivity γ_τ ramps up linearly over the first $R = 3$ rounds, following $\gamma_\tau = \gamma \cdot \min\left\{\frac{\tau+1}{R}, 1\right\}$.

Agent-specific withdrawal thresholds θ_i are drawn from a truncated normal distribution $\mathcal{N}(\bar{\theta}, \sigma_\theta^2) \cap (0, 1)$, with a mean of $\bar{\theta} = 0.56$ and a standard deviation of $\sigma_\theta = 0.18$. Agents who receive neutral signals (y_M) redeposit deterministically whenever their updated beliefs fall below their individual thresholds, that is, when $b_{i\tau} < \theta_i$. The simulation proceeds until it reaches either two consecutive rounds with no belief changes or a maximum of 30 rounds, whichever condition is met first.

Mental Models

Table A.3: Example Explanations by Factor (Withdrawal = Yes)

Factor	**Human Examples**	**Gemma Examples**
FDIC Insurance	It's not insured and there seem to be problems/instability with banks.	I'm worried about losing my savings, especially the uninsured portion. This tweet, even if it's alarmist, is making me nervous.
Risk and Safety	Not taking any risks with my money.	I'm close to retirement and can't risk it; I'll take my money out now.
Trust and Doubt	I'd rather have my money in my pocket than trust anybody else, even a bank's.	I can't afford to lose what little I have. That tweet scared me and I don't trust the bank's email.
Information/Evidence	News article and a tweet are no coincidence; even if I don't withdraw all my money, I'd take out a substantial amount.	That tweet scared me, even though the news says things are fine.
Financial Constraints	I live paycheck to paycheck; I can't afford to lose what little I have.	I don't understand all that talk about the markets, but I can't afford to lose my money. Better safe than sorry.
Social/Panic	If it's spreading beyond a couple of local banks, it's not a coincidence—rumors become reality when enough people panic...	Seeing everyone freaking out online about another bank is scary; I don't want to risk losing my savings.
Institutional Credibility	I can't trust this—don't you know better than to trust the government? I'd take enough money out...	I don't trust it when the bank says they're fine; too many people are saying bad things.
Personal Experience	I've had problems with my bank in the past and don't trust them on a normal day; that tweet would make me withdraw for peace of mind.	I'm young and don't have a lot of financial experience. That tweet scared me, and I don't want to risk losing my savings...

Notes: Examples are drawn from free-text responses using the shared taxonomy. Some responses are truncated for readability.

Table A.4: Representative Explanations by Factor (Stay = No Withdrawal)

Factor	**Human examples**	**Gemma examples**
FDIC Insurance	FDIC will cover it even if it's uninsured; they almost always make it right.	I trust the deposit insurance framework and don't see a reason to panic.
Risk and Safety	I think money in the bank is safe. It would take something catastrophic for my funds to be in danger.	I've been with my bank for years and they say funds are safe.
Trust and Doubt	Information found online isn't really reliable; I'd wait for a more trustworthy source.	My bank contact says everything is fine, and I don't trust random tweets over a trusted source.
Information/Evidence	There was no evidence in the post; if there were such an emergency, it would be headline news.	I prefer verified statements to rumors on social media.
Financial Constraints	I believe banks are secure for holding money; I don't have enough on deposit to worry about.	The President said the banks are secure and I trust that; I need my money for bills.
Social/Panic	One tweet about getting money out—when it's not even my bank or state—doesn't justify panic.	The email from my bank sounds reassuring; I don't want to panic based on a random tweet.
Institutional Credibility	One tweet about withdrawals from a different bank or state doesn't persuade me.	The email from my bank is reassuring; I won't panic based on a tweet.
Personal Experience	I worked at a bank part-time during 9/11 and I vividly remember the mad dashes.. I thought it was ridiculous then and I do today, too.	The Fed Chair said the banking system is strong, and I trust that; I've been through ups and downs and panicking never helped.

Notes: Examples are drawn from free-text responses using the shared taxonomy. Some responses truncated for readability.