Self-fulfilling Banking Crises:

An Overview

Todd Keister

Lectures delivered at Washington University
February 2007

Banking crises

- Banking crisis: an event where many banks fail or are bailed out by government
  - occur frequently and in a wide variety of countries
  - 86 major episodes worldwide between 1974 and 1997
  - costly to depositors and/or taxpayers
  - often have important macroeconomic consequences
- Many things are happening during these banking crises

Bank runs

- Banking crises sometimes include a “run”
  - sudden rush to withdraw by many depositors
- Runs occurred frequently in U.S. before 1933
  - “major” events in 1890, 1893, 1899, 1901, 1903, 1907, 1932-3

American Union Bank, New York City, 1933

- Sometimes a run involves an individual bank(s)
  - not part of a general crisis

Abacus Federal Savings Bank, New York City
April 22, 2003

- Our focus: system-wide runs
  - more important from a macroeconomic perspective

- Recent system-wide runs: Argentina in 2001, Russia in 2004

Total deposits in banking system in Argentina in 2001

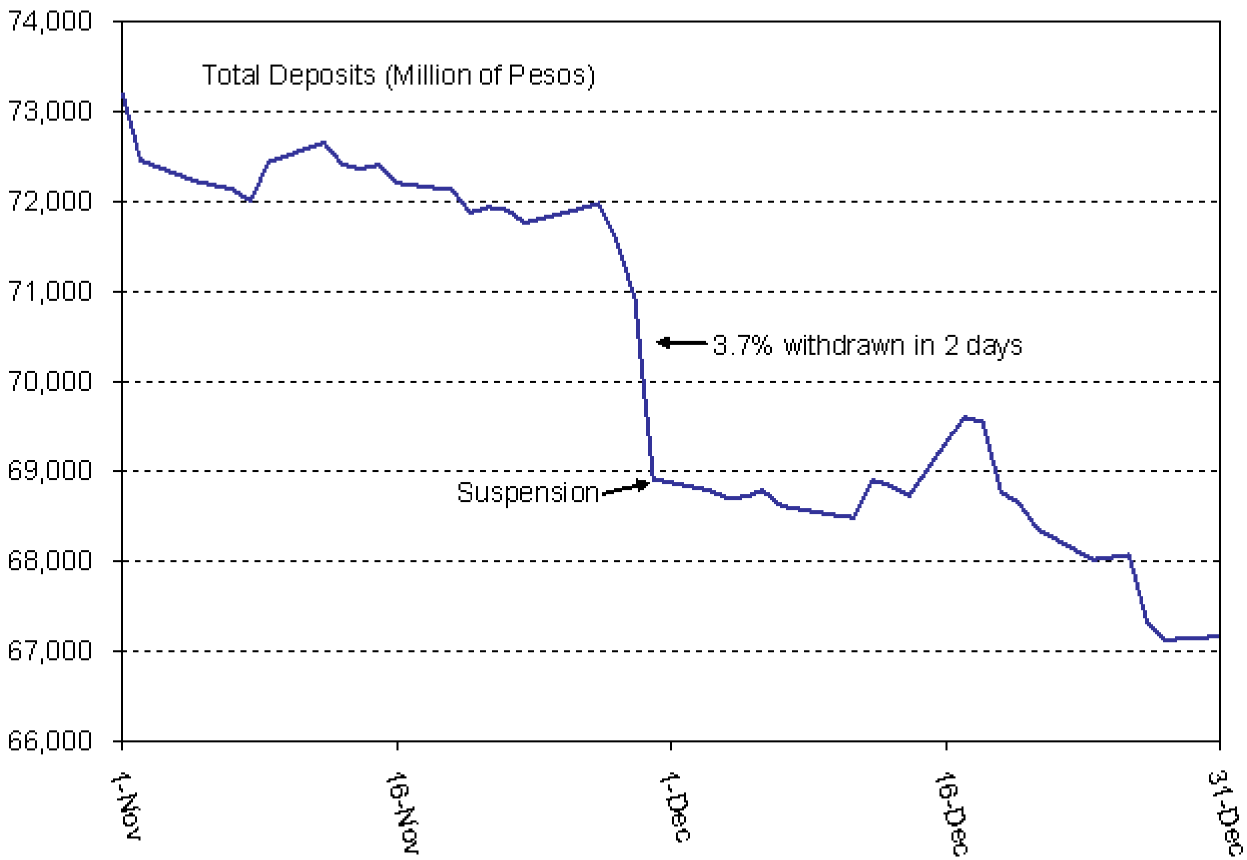

Q: What causes these runs?

Two competing views

- "Fundamental" causes
  - economic downturn makes banks insolvent (cannot afford to repay all depositors)
  - depositors learn this, rush to withdraw
  - those who are late to withdraw lose money
- In this view, the run is a *symptom* of the overall crisis
  - banks are already insolvent; losses have occurred
  - run is the final event before closure/bailout

Other view

- Self-fulfilling beliefs
  - depositors worry that bank will fail

  $\Rightarrow$ rush to withdraw
  - sudden large withdrawal demand causes bank to fail

- In this view, the run *causes* or *deepens* the crisis
  - banks were solvent (maybe some losses, but fundamentally OK)
  - run forces them to liquidate assets, *generates* losses

- Importance of self-fulfilling beliefs often stressed by observers

  J.P. Morgan during the crisis of 1907: “If the people will keep their money in the banks everything will be all right.”

- Also deeply rooted in popular culture

  – *It’s a Wonderful Life* (1946), *Mary Poppins* (1964)

- Important to understand the underlying cause of runs
  - if fundamental, perhaps no reason to prevent runs
  - if self-fulfilling, crisis and resulting downturn are not inevitable; perhaps can be prevented

- Policy implications: government-provided deposit insurance
  - common; often effective at preventing runs
  - generates a moral hazard problem
  - costly, as seen in Savings & Loan crisis ($125 billion)
  - benefits depend on what (potentially) causes runs

- Determining the underlying cause of a run is very difficult
  - real world crises are complex, with many confounding factors
  - answering counterfactual questions is difficult/impossible

  "What would have happened in Argentina in 2001 if people had left their money in the banking system?"

- A research agenda: Are self-fulfilling bank runs plausible?
  - can we write a (rational, optimizing) model in which they occur?
  - if not, strong evidence in favor of the fundamental-cause hypothesis
  - if so, we should take the possibility seriously (example: be cautious about deposit insurance reform)

Outline of lectures

- The Diamond and Dybvig (JPE, 1983) model
  - the seminal model of self-fulfilling runs
  - generates a bank run equilibrium, but ...
- Subsequent criticism
  - is the run equilibrium robust?
  - does it depend on special assumptions?
- Overview of some recent work
  - no consensus answer

Constructing a model

- Potential problem arises from combination of **fractional reserve banking** and **demand deposit contracts**

- Effect: bank's short-term liabilities are larger than assets on hand
  - maturity mismatch

Q: Why do banks operate this way?

- Want to write down an environment where:
  - there is a role for banks, and
  - banks choose to have a maturity mismatch

Diamond and Dybvig (JPE, 1983)

Outline

- The physical environment
- The efficient allocation
- How to decentralize this allocation
  - banks
- Equilibrium

The environment

- Single commodity
- 3 time periods, $t = 0, 1, 2$
- Continuum of (ex ante) identical depositors (measure 1)
- Endowment: each agent has 1 at $t = 0$
  - nothing later
- Utility: $u(c_1 + \theta c_2)$

$$\text{where } \theta = \left\{ \begin{array}{c} 0 \\ 1 \end{array} \right\} \text{ if depositor is } \left\{ \begin{array}{l} \text{impatient} \\ \text{patient} \end{array} \right\}$$

- Type (patient or impatient) is revealed only at $t = 1$
  - type is private information

- Assume: each agent is impatient with probability π
  - fraction π of agents will be impatient (non-stochastic)

- Two technologies

	Return at $t = 1$	Return at $t = 2$
storage	1	1
investment	$1 - \tau$	$R > 1$

where $\tau \geq 0$ is a *liquidation cost*

The (full information) efficient allocation

- Suppose a social planner could observe types and assign allocations
  - find the symmetric allocation that maximizes expected utility at $t = 0$

- Clearly impatient depositors will consume only at $t = 1$; Amount: c_E
  - patient depositors will consume only at $t = 2$; Amount: c_L

- Expected utility

$$\pi u(c_E) + (1 - \pi) u(c_L)$$

- Feasibility:

$$\begin{aligned} \pi c_E &= 1 - i \\ (1 - \pi) c_L &= Ri \end{aligned}$$

- Planner's problem:

$$\max \ \pi u(c_E) + (1-\pi)\, u(c_L)$$

subject to

$$\pi c_E + (1-\pi)\frac{c_L}{R} = 1$$

and

$$c_E \geq 0, c_L \geq 0$$

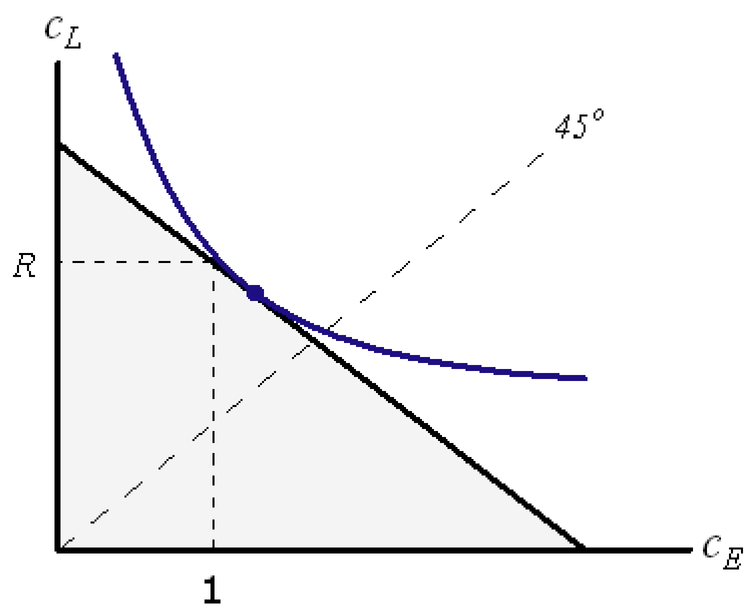

- First-order condition:

$$u'(c_E) = Ru'(c_L)$$

- Solution: $\left(c_E^*, c_L^*\right)$ (note: $c_L^* > c_E^*$ because $R > 1$)

- Assume

$$c_E^* > (1 - i^*) + (1 - \tau)\, i^* = 1 - \tau i^*$$

  – optimal allocation provides "liquidity insurance"

- Holds as long as risk aversion is high enough

  – if $\tau = 0$, holds if and only if coefficient of RRA is > 1

- Notice: no liquidation occurs in the efficient allocation

Implementing the efficient allocation

- Suppose an institution called a "bank" arises

- Offers the following *demand deposit contract*

- In exchange for your endowment at $t = 0$, you have the right to

  – a fixed payment c_E^* at $t = 1$ (as long as the bank has funds)

  or: a pro-rata share of the bank's assets at $t = 2$ (hopefully c_L^*)

- Resembles a demand deposit contract in that agent chooses when to withdraw her money

- Suppose all agents accept the contract
  - bank collects funds, sets $i = i^*$

- Agents then play a *withdrawal game*
  - each observes her own type and then decides whether to withdraw at $t = 1$ or at $t = 2$

- Strategy: choose withdrawal date as function of type θ
  - impatient depositors will always withdraw early (dominant strategy)
  - only need to consider what agent will do when patient $s_i \in \{1, 2\}$

- Let s denote the profile of all withdrawal strategies; $s : [0, 1] \longrightarrow \{1, 2\}$
  - s_{-i} is the profile of all strategies except agent i

- Look for Nash equilibria of the withdrawal game
  - focus on symmetric, pure strategy equilibria

- Suppose a depositor believes that all others will wait until $t = 2$ when patient (that is $s_{-i} = 2$)
  - if she withdraws at $t = 1$, she receives c_E^*
  - if she withdraws at $t = 2$, she receives c_L^* $(> c_E^*)$

  $\Rightarrow$ optimal action is to wait when patient

- Repeating:
  - if she withdraws at $t = 1$, she receives c_E^*
  - if she withdraws at $t = 2$, she receives c_L^* $(> c_E^*)$

		all others (s_{-i}) 1	all others (s_{-i}) 2
agent i	$s_i = 1$	–	c_E^*
	$s_i = 2$	–	c_L^*

$\Rightarrow$ (Full information) efficient equilibrium is a Nash equilibrium

- demand deposit contracts implement the efficient allocation (even though types are private information!)

An aside

- Recall: We said: "Suppose a bank were to ..."
  - then the efficient allocation can be achieved

- Would profit-maximizing banks offer the contract $\left(c_E^*, c_L^*\right)$?

- Competition: suppose banks announce contracts, depositors are free to choose any bank

- Equilibrium contract must maximize depositors' utility
  - otherwise someone could change contract slightly, gain all deposits
  - Bertrand competition

Bank runs

Q: Are there other equilibria of the withdrawal game?

- Suppose a patient depositor believes that everyone else will try to withdraw at $t = 1$ (that is, $s_{-i} = 1$)

- Some depositors will receive c_E^*, but ...
  - after paying a fraction

$$\chi \equiv \frac{1 - \tau i^*}{c_E^*} (< 1)$$

    bank's assets are depleted

  - remaining depositors receive nothing

- If agent i waits until $t = 2$, she receives 0 for sure

- Looking at the payoffs:

		all others (s_{-i}) 1	2
agent i	$s_i = 1$ |	c_E^* or 0	c_E^*
	$s_i = 2$ |	0	c_L^*

- There is a Nash equilibrium in which all depositors try to withdraw at $t = 1$
  - resembles a bank run

$\Rightarrow$ This is the Diamond-Dybvig theory of bank runs

A recap

- Banks are illiquid ...
  - short-term liabilities $>$ short-term assets
- ... because it solves a private information problem
- Demand deposit contracts and fractional reserve banking are a way to implement the efficient allocation
- But ... they also open the door to self-fulfilling runs

Possible criticisms

- Recall our original question:

  Can we write down a (rational, optimizing) model in which bank runs emerge as an equilibrium outcome?

- Is the Diamond-Dybvig model a satisfactory answer?

Three possible criticisms:

(1) If agents have rational expectations, they should know the run will occur and not deposit in bank

(2) Contingent contracts could eliminate the run equilibrium

(3) Suspension of payments could eliminate the run equilibrium

(1) If agents have rational expectations, they should know the run will occur

   - should not deposit with this bank

   - choose autarky or some other arrangement

- U.S. history: panics every 3-5 years could not be entirely unexpected

- Need to look at the "overall" game

   - includes the choice of banking contract, deposit decision

Q: What should the bank do if it expects depositors to run?

- Can the bank offer a deposit contract that convinces depositors not to run?

  - what would be required?

- Need $c_E \leq 1 - \tau i$

  - if everyone attempts to withdraw at $t = 1$, all will be served

  - in this case, a patient depositor will prefer to wait

  - "run proof" contract (Cooper and Ross, JME 1998)

- Can be satisfied by lowering c_E and/or investment

- Rewriting to the bank's constraints:

$$\begin{aligned} \pi c_E &= 1 - i - r \\ (1-\pi)\, c_L &= Ri + r \end{aligned}$$

- r is "excess reserves"; inefficient, but increase $1 - \tau i$

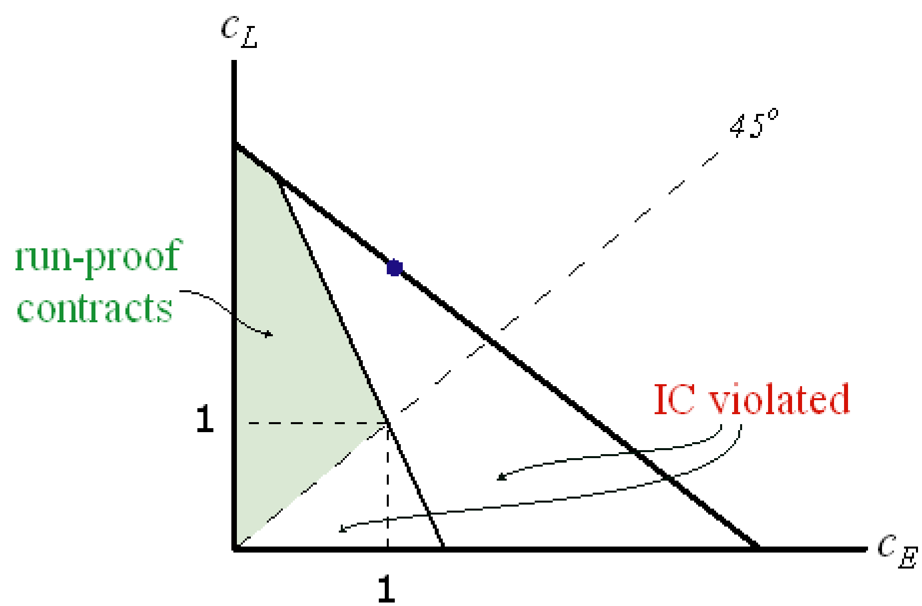

The "overall" banking game:

- At $t = 0$, bank sets $c = (c_E, i, r)$, anticipating the actions of depositors
  - chooses a point on previous graph
- At $t = 1$, depositors learn types and play the withdrawal game
  - withdrawal game is a proper subgame
  - bank moves first, can influence actions of depositors
  - Stackelberg leader
- What is an equilibrium of this overall game?

- For each possible c, bank forecasts depositors' actions in withdrawal game
  - chooses c to maximize expected utility at $t = 0$ conditional on this forecast

Formally:

- A (subgame perfect) Nash equilibrium of the overall banking game is a contract $\widehat{c}$ and a function $\widehat{s}(c)$ such that

(i) for every c, $\widehat{s}(c)$ is an equilibrium of the withdrawal game, and

(ii) $\widehat{c}$ maximizes $V(c, \widehat{s}(c))$

Q: What are the (symmetric) Nash equilibria of this overall game?

An equilibrium with no bank run:

- Suppose $\widehat{s}(c) = 2$ for all relevant c
  - patient depositors never run
  - then c^* clearly maximizes expected utility
  - c^* together with $\widehat{s}(c) = 2$ is an equilibrium

$\Rightarrow$ The efficient allocation is an equilibrium of the overall game

- Is there also a bank run equilibrium?

Looking for a bank run equilibrium:

- Suppose $\widehat{s}\,(c^*) = 1$
  - depositors will run if bank chooses c^*

- Then the bank will not offer c^*
  - it will choose the best run-proof contract $\widetilde{c}$
  - by definition, depositors will not run at $\widetilde{c}$

- There cannot be an equilibrium with $\widehat{s}\,(c) = 1$ for the chosen c

$\Rightarrow$ a bank run cannot occur in a Nash equilibrium of the overall game

An aside:

- There *is* an inefficient equilibrium in this game
  - $\widehat{s}(c) = 1$ for all c that are not run proof
  - bank chooses best run-proof contract $\widetilde{c}$
- Bank runs need not occur in order to generate welfare losses
  - possibility of a run can distort the banking contract
- There could be a role for deposit insurance even if runs are never observed
- However, we do observe runs, so ...

- Does this sink the Diamond-Dybvig theory of bank runs?
  - what is missing?
- We have shown that bank runs cannot occur with perfect foresight
  - not very surprising
  - need some uncertainty
- Need for depositors (and the bank) to be optimistic at $t = 0$
  - then a "bad" shock occurs at $t = 1$
- Can be captured in a *correlated equilibrium*

- Introduce an extrinsic state $\sigma \in [0, 1]$; "sunspots"
  - realized at $t = 1$
  - observed by depositors, but not by bank

Formally:

- A correlated equilibrium of the overall banking game is a contract $\widehat{c}$ and a function $\widehat{s}(c, \sigma)$ such that

(*i*) for every c and s, $\widehat{s}(c, \sigma)$ is an equilibrium of the withdrawal game, and

(*ii*) $\widehat{c}$ maximizes $\int V(c, \widehat{s}(c, \sigma)) \, d\sigma$

- Story: everyone knows a run will occur if enough spots appear on the sun (or if stock market falls more than x%)
  - sunspots are exogenous event; contain no information
  - $\sigma =$ quantity of sunspots; run if $\sigma > \widehat{\sigma}$

- at $t = 0$, contract is set based on this prior belief

- At $t = 1$, σ is realized
  - if low – good; no run; efficient allocation obtains
  - if high – bad; depositors run

Result: (Cooper and Ross, 1998, and others)

There exists a correlated equilibrium of the overall game in which a bank run occurs with positive probability if and only if there exists a bank run equilibrium of the withdrawal game at c^*.

- Intuition: if run is unlikely, bank will choose contract close to c^*
  - then a run is an equilibrium of the withdrawal game

- Bottom line: theory is completely consistent with rational expectations

- From now on, we focus only on the withdrawal game and the efficient allocation c^*

(2) Contingent contracts:

- Suppose bank can make payments contingent on early withdrawal demand
  - if fraction θ want to withdraw early, calculate $c_E^*(\theta)$ and $c_L^*(\theta)$
  - as in Champ, Smith, and Williamson (1997)

- Easy to show:

$$c_L^*(\theta) \geq c_E^*(\theta) \quad \text{for any } \theta \in (0,1)$$

  - waiting is dominant strategy for a patient depositor

- Unique equilibrium. Bank runs do not occur.

- Is this really a solution to the bank-run problem?

- Approach requires bank to ask all depositors if they want to withdraw early before making *any* early payments
  - does not seem like a demand deposit contract

- Wallace (1988): *Sequential service* is an essential feature of banking
  - depositors arrive at bank sequentially, must be paid as they arrive
  - this is a feature of the environment

- Result: information about θ is only gradually revealed to bank over time

- First depositor arrives, bank must choose c_E
  - has no information about total withdrawal demand

  $\Rightarrow$ early payments cannot depend on θ

$\Rightarrow$ This type of contingent contract cannot be used

  - point was also made in Diamond and Dybvig (1983)

- Important: sequential service is a restriction on what is feasible in the environment
  - not an (ad hoc) restriction on banks

An aside:

- Some authors have argued that banks are not essential in the Diamond-Dybvig environment
  - a mutual fund with market-traded shares can also implement the efficient allocation (Jacklin, 1987)
  - no possibility of a run

- The sequential service constraint implies that such markets *cannot* operate in the (true) Diamond-Dybvig environment

- Other authors study the interaction of bank and markets
  - always need to be careful about what markets are feasible and whether banks are useful

(3) Suspension of payments

- Payment to first depositor cannot depend on θ, but what about later depositors?
  - bank eventually learns a run is underway
  - could use a contingent contract at that point

- Suppose bank modifies the deposit contract with a *suspension clause*
  - after giving π depositors c_E^* at $t = 1$, bank will close doors
  - no further withdrawals until $t = 2$;
  - "suspends convertibility of deposits into currency"

- A patient agent knows only π withdrawals will occur at $t = 1$
  - if she waits until $t = 2$, bank will be able to pay $c_L^* \left(> c_E^*\right)$
  - regardless of what other agents try to do

$\Rightarrow$ waiting is a dominant strategy for patient agents

- There is a unique equilibrium
  - only impatient agents withdraw at $t = 1$; bank runs cannot occur

- Note: suspension never occurs in equilibrium
  - the threat of suspension convinces depositors not to run

- This criticism is less easy to overcome

- One approach: assume banks cannot suspend payments
  - required to pay c_E to all depositors until assets are gone
  - justify using stories about moral hazard among bankers
  - widely used (Cooper and Ross 1998, Chang and Velasco 2000, many others)

- However, this approach is not very satisfactory
  - depends on an *ad hoc* restriction on contracts
  - open question: could one write a *model* of this?

Other possible approaches:

- Assume π is stochastic
  - bank does not know precise point at which to suspend payments
  - uncertainty about π must be large
  - approach taken by Diamond-Dybvig, Wallace (QR, 1988 and 1990), Green and Lin (JET, 2003), Peck and Shell (JPE, 2003)

- Results:
  - optimal deposit contract becomes complex; no longer resembles a demand deposit
  - run equilibrium may or may not exist depending on informational assumptions

Or:

- Assume bank cannot *commit* to suspend payments

- Recall the timing:

  - π agents withdraw, then bank realizes a run is underway

  - some agents still in line are (truly) impatient

  - suspending payments means they receive nothing

- Without commitment, how would bank (or perhaps govt.) respond to a run?

  - ask if the response gives agents an incentive to run

Overview of recent work

(i) Banking policy without commitment

- study *ex post* optimal response to a run
- derive conditions under which a run equilibrium exists
- based on Ennis and Keister (2006)

(ii) Banking with aggregate uncertainty

- study two different specifications of sequential service (Peck and Shell, JPE 2003; Green and Lin, JET 2003)
- derive the efficient allocation under each
- ask if a bank run equilibrium exists

Banking policy without commitment

Q: Without commitment, is the threat to suspend payments credible?

- After π withdrawals at $t = 1$, banking authority learns a run is underway (sequential service)
  - some of the *past* withdrawals were by patient depositors
  - some depositors still in line are (truly) impatient
- Suspending payments implies giving nothing to these depositors

$\Rightarrow$ Banking authority might not want to suspend right away

- threat that eliminates the run equilibrium might not be credible
- patient depositors, recognizing this, might choose to run

Example: Suspension in Argentina

Total deposits in banking system, January 2001 – June 2002

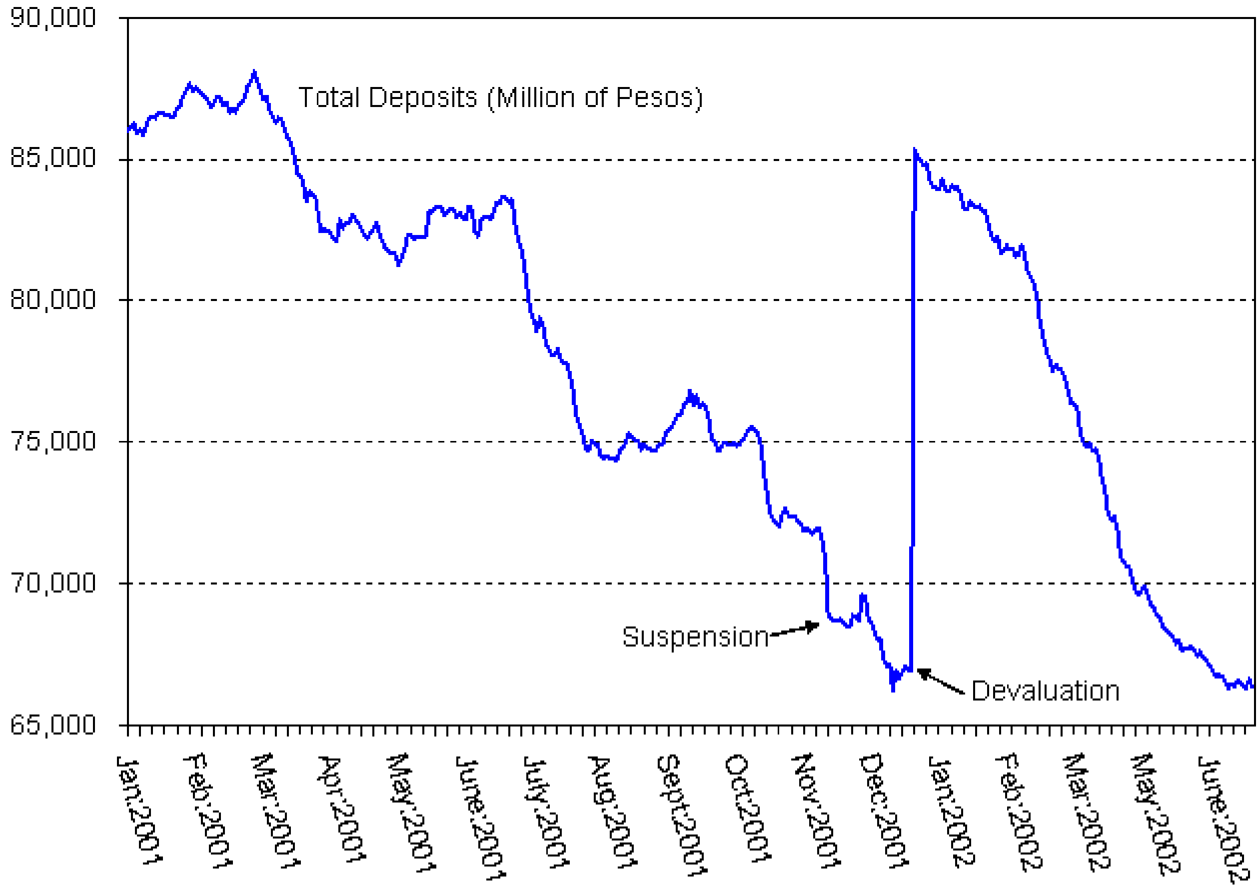

Point: Implementing a suspension may be difficult/undesirable *ex post*

Credible simple suspension policies

- Suppose a run is underway and π depositors have already received c_E^*
  - $1 - \pi$ remain in line; a fraction π of these are impatient

- In a *simple suspension scheme*, banking authority must either

  (i) continue to pay c_E^*, or

  (ii) suspend payments until $t = 2$

- Ask: at what point would the bank choose to suspend?

- The credible suspension point π_s solves

$$\max_{\pi \leq \pi_s \leq 1} \; (\pi_s - \pi)\, u\left(c_E^*\right) + (1 - \pi_s)(1 - \pi)\, u\left(c_L\left(\pi_s\right)\right)$$

subject to

$$(1 - \pi_s)(1 - \pi)\, c_L = R\left(i^* - \frac{(\pi_s - \pi)\, c_E^*}{1 - \tau}\right)$$

  – additional payments at $t = 1$ must come from liquidating investment

- Solution: π_s^*

- A patient depositor who expects a run followed by a suspension gets:
  - $c_L(\pi_s^*)$ if she waits
  - either c_E^* or $c_L(\pi_s^*)$ if she runs

- If $c_E^* > c_L(\pi_s^*)$ she would choose to run
  - happens if π_s^* is large enough
  - suspension of convertibility does *not* eliminate the run equilibrium

- Proposition: The run equilibrium exists under the credible suspension policy if and only if

$$\pi > \frac{1-\gamma}{\gamma}\left(\frac{R}{1-\tau}-1\right) \quad \text{when } u(c)=\frac{c^{1-\gamma}}{1-\gamma}$$

- Repeating the condition:

$$\pi > \frac{\gamma}{1-\gamma}\left(\frac{R}{1-\tau} - 1\right)$$

will hold if:

- π is large
  - many of the depositors still in line are impatient, or
- γ is small
  - as in Diamond-Dybvig, bank runs are a problem if depositors are sufficiently risk averse

Court intervention

- One striking feature of the crisis in Argentina was the *ex post* involvement of the courts
  - depositors could claim "special need" to withdraw full deposits
  - illness and hospitalization were common claims
- Almost 200,000 cases were resolved in favor of depositors
- Over 10% of deposits paid out this way after suspension was declared

- Adding courts to the model
  - suppose that by paying a cost α, a "court" can determine a depositor's true type
  - set $\alpha = 0$ for now, but can only intervene once a run is underway
  - court can raise welfare by directing additional payments to impatient depositors

- Court can choose any continuation payments
  - impatient depositors served after π may receive less than c_E^*
  - sometimes called a *partial* suspension

- Court's problem:

$$\max_{c_{ER},c_{LR}} \; (1-\pi)\left[\pi\frac{1}{\gamma}(c_{ER})^{\gamma}+(1-\pi)\frac{1}{\gamma}(c_{LR})^{\gamma}\right]$$

subject to

$$\begin{aligned}(1-\pi)\,\pi c_{ER} &\leq (1-\tau)\,i^{*}\\ (1-\pi)^{2}c_{LR} &= R\left[i^{*}-\frac{(1-\pi)\,\pi c_{ER}}{(1-\tau)}\right]\end{aligned}$$

- Solution: $\left(c^{*}_{ER}, c^{*}_{LR}\right)$

  – looks like planner's problem in Diamond-Dybvig

- Note: this solution is the *first-best* continuation allocation

Results:

- Proposition: There exist parameter values for which $c_E^* > c_{LR}^*$ holds
  - there is an equilibrium where depositors run at the beginning of $t = 1$
  - courts intervene; implement the efficient continuation allocation; but ...
  - all depositors withdrawing after the suspension occurs receive less than those who withdrew early
- Happens when
  - liquidation costs are high enough (τ is large), or
  - depositors are sufficiently risk averse (γ is small)

- Instead of courts, banking regulator could reschedule payments
  - with no ability to screen types

- Analysis is exactly as before *if* the payment rescheduling halts the run

- New issue: might patient depositors continue running?
  - there cannot be an equilibrium where they continue running for sure, but ...
  - can introduce another source of (extrinsic) uncertainty
  - there are equilibria where the run continues with positive probability, payments are rescheduled again, etc.

Bottom line:

- Recall: self-fulfilling runs can only occur if banks do not suspend payments at π
- Why don't they?
  - one answer: suspending payments is *ex post* inefficient, **and**
  - committing to suspend is difficult (Argentina in 2001)
- If banking authorities cannot commit to a strict suspension policy
  - response when faced with a run will be lenient
  - further compromises solvency of the banking system
  - can justify the *initial* decision of depositors to run

Banking with aggregate uncertainty

- Another reason banks might not suspend payments:
  - π is random
  - do not know exactly when to suspend
- Assume: each depositor observes own type; no one observes π
  - different from Champ-Smith-Williamson
- When π is random the efficient allocation is different
  - sequential service: π is unknown when first depositor withdraws

Proceed as before

- Specify the environment

- Find the efficient allocation c^*

- Ask if banks can implement c^* as *an* equilibrium

- Ask if a run equilibrium exists at c^*

  – if so, then a run equilibrium exists in the overall game (as before)

Peck and Shell (JPE, 2003)

- Environment is like Diamond-Dybvig except:
  - finite number of depositors (I)
  - π is random (with full support)
  - $\tau = 0$ (no liquidation cost; no portfolio choice)

- With real uncertainty, the planner's problem is more complex
  - sequential service restricts the flow of information
  - planner learns π gradually

Timing:

- First impatient depositor arrives at bank
  - planner only knows there is at least 1 impatient depositor
  - pays some amount c_E^1
- Second impatient depositor arrives
  - planner knows there are at least 2 impatient depositors
  - pays c_E^2, ... and so on
- After all impatient depositors have arrived:
  - remaining depositors share resources at $t = 2$

- Let $p_i =$ Prob(at least i are impatient $|$ at least $i-1$ are impatient)

- Let $V_i(y) =$ value of having y remaining when i^{th} impatient depositor arrives

- Planner's problem:

$$V_i(y_{i-1}) = \max_{\{c_E^i\}} \left\{ \begin{array}{c} u\left(c_E^i\right) + p_{i+1} V_{i+1}\left(y_{i-1} - c_E^i\right) + \\ \\ (1 - p_{i+1})(I - i)\, u\left(\frac{y_{i-1} - c_E^i}{I - i}\right) \end{array} \right\}$$

with

$$V_I(y_{I-1}) = u(y_{I-1}) \qquad \text{and} \qquad y_0 = I$$

- Solution: $c_E^* = \left\{c_E^{1*}, ..., c_E^{I*}\right\}$

Decentralizing c^*:

- Banking contract: demand deposit, but
  - how much you receive depends on when you withdraw

- Withdrawal game: as before, except:
  - value of withdrawing is random
  - need to compare expected utilities

- Note: no scope for suspending payments
  - efficient allocation requires that c^* be followed

Result: (Peck and Shell, 2003)

- For some parameter values, a run equilibrium exists

Bottom line:

- Why don't banks eliminate the run equilibrium by suspending payments?
  - another answer: doing so would hurt some impatient agents in some states
  - it would be *ex ante* inefficient
  - requires "significant" uncertainty about π

Green and Lin (JET, 2003)

- Environment is slightly different from Peck-Shell
- Argue that depositors should be able to observe c_E^i before deciding
- Capture this feature by assuming:
  - all depositors go to bank at $t = 1$
  - observe their position in the "line", then decide when to withdraw
  - value of waiting is random, but withdrawing early offers a known amount

- Since depositors know their position in line, each faces a different problem

- Focus on the last depositor in line
  - can show: she always receives more by waiting
  - reason: no one left to "run"
  - waiting when patient is a dominant strategy for her

- Proceed by backward induction
  - calculations are messy, but result is striking

Result: (Green and Lin, 2003)

- The efficient allocation c^* is the *unique* equilibrium of the withdrawal game

Bottom line:

- When depositors have more information, the run equilibrium may disappear
  - even if bank never suspends payments

- Some people interpret this as evidence for the "fundamental cause" hypothesis

- Open question: how robust is their result?

Final remarks

- Are self-fulfilling bank runs theoretically plausible?
  - open to debate

- One view: the question of commitment is likely to be central
  - the "right" policies by banks/govt. may be able to prevent self-fulfilling runs
  - but ... will these policies be followed?

- Room for more research