Lecture 20: Intro to decision theory & game theory

Complex Systems 530

Credit: Many of these slides have been borrowed directly from lectures by Michael Hayashi and Lynette Shaw!

How smart are your agents?

Reflexive Agents
simple, static rules

Goal-Based Agents
rules adjust according to being in goal state

Utility-Based Agents
rules attempt to maximize utility function(s)

Adaptive Agents
rules update based on experience

- ***Cognitive Complexity*** **+**

Game theory

- Game theory - motivated by the realization that the study of strategically interdependent behavior can be greatly enhanced via analysis of mathematical models of conflict and cooperation between “rational” decision-makers

- First got going as a field in 1940s after publication of Theory of Games and Economics work by von Neumann and Morgenstern

Introduction

- Applied to a wide range of areas
  - Social sciences (economics, sociology, political science)
  - Biology (genetics, species)
  - Computer science and logic
- Basic idea is that if we can conceptualize the interdependencies of individuals in a system as a game, will be able to “solve” for outcomes (for individuals all the way to population levels)

Decision theory

- Sort of a one-player version of game theory, where each person decides an action based on their preferences and the expected outcome of their actions (but no considering of other individuals/ players involved)

Decision theory

- **Actions**: The set of things an individual can do. e.g. video games, nap, run simulations.

- **Outcomes**: The results of each action
  - Video games → entertainment
  - Nap → rest
  - Run simulations → work

- **Preferences**: An ordering that specifies how an individual ranks the outcomes.
  entertainment > rest > work

Preferences

- For a preference order to be **rational**, it must be *complete* and *transitive*.
- **Complete**: for every pair of outcomes, one is preferred over the other (or they can be indifferent). Formally, for every a and b, $a > b$, $b < a$, or $a = b$. (One can consider strict $>$ or weak $\succcurlyeq$ preference)
- **Transitive**: For any three outcomes, a, b, and c, if a is preferred to b, and b is preferred to c, then a must be preferred to c. Formally, $a > b$ and $b > c$ implies $a > c$.

Rationality

- Completeness and transitivity guarantee that a person will be able to identify the best alternative out of their available options
- A rational actor in the economic sense always picks the most preferred alternative
- Note that a rational choice $\neq$ good choice!

Preferences

- Preferences are often described using a **utility function** or **payoff function**, which assigns a number/value to each outcome—the ordering is then assessed based on the utility function value
- Individuals then attempt to choose their actions to maximize their utility

Game theory

- Most decisions aren't made in isolation – it's important to know what somebody else might do. Game theory extends decision theory to problems where other people are a factor.

Game theory

- **Game**
  - Circumstances where results depends on the actions of 2 or more individuals (players)
  - Outcomes (payoff structures) are knowable and pre-defined
- **Players**
  - Possess choices (strategies) they can play
  - Seek to maximize their own utility/payoff (self-interest) and have the information and cognitive capacity to do so (rationality)
  - Typically everybody has common knowledge

Game theory

- **Players**: the actors making decisions
- **Strategies**: sets of choices specified for each player—these may or may not be the same across players!
- **Strategy profile**: a set containing one strategy chosen by each player
- **Payoffs**: a numerical representation of the costs and benefits of each strategy profile to each player

Common assumptions for games

- **Rationality**: Each player picks the action that gives the highest payoff given what they believe the other player might do. (Players always play best responses)

- **Complete information**: Each player knows the game, all of the payoffs, and all of the actions available to every player.

- **Common knowledge:** Each player knows that the other players are rational and have complete information.

Common assumptions for games

- In practice, at least one of these assumptions is often violated. The basic theory shown here can be extended to deal with some deviations from rationality (bounded rationality, evolutionary game theory) and incomplete information (Bayesian games, and others).

Game variants

- Games come in a wide number of varieties:
  - Non-cooperative vs. cooperative
  - Zero-sum vs. non-zero-sum
  - One shot vs. Iterated
  - Symmetric vs. Non-symmetric
  - Simultaneous vs. Sequential
    (related: Normal vs. Extensive forms)
  - Two vs. Many player

Solutions to Games

- Anticipating the outcome of a games is often oriented toward analytically solving for the “stable” configuration of choices individuals can make

- Specifically, oriented toward identifying the **Nash equilibria** of a game:

    - Given that all players know each others’ equilibrium strategies, no player can benefit from changing their own strategy while the other players’ strategies remain unchanged

Nash Equilibria

- **Nash's Existence Theorem**: if mixed strategies (where a player chooses probabilities of using various pure strategies) are allowed, then every game with a finite number of players in which each player can choose from finitely many pure strategies has at least one Nash equilibrium.
- What if infinite set of choices? Can have issues if non-compact set:
  - Two players simultaneously name a number and the player naming the larger number wins.
  - Two players choose a real number strictly less than X and the winner is whoever has the biggest number

Prisoner’s Dilemma

		Prisoner A	
		Silent	Betray
Prisoner B	Silent	-1 -1	0 -3
	Betray	-3 0	-2 -2

Prisoner’s Dilemma

Nash equilibria?

Prisoner’s Dilemma

		Prisoner A	
		Silent	Betray
Prisoner B	Silent	-1 -1	0 -3
	Betray	-3 0	-2 -2

Prisoner’s Dilemma

Pure coordination game

	L	R
L	10,10	0,0
R	0,0	10,10

E.g.—suppose walking and don't want to bump into the person walking the other way—want to both swerve same direction (e.g. if both swerve to their own right, will miss each other, but if one swerves to their right and the other to their left they will bump into one another!)

Pure coordination game

	L	R
L	10,10	0,0
R	0,0	10,10

Multiple Nash equilibria

Other coordination games

Assurance game

	Party	Home
Party	10,10	0,0
Home	0,0	5,5

BoS

	Bach	Stravinsky
Bach	10,5	0,0
Stravinsky	0,0	5,10

Stag hunt

	Stag	Hare
Stag	10,10	0,5
Hare	5,0	5,5

(Or sometimes, 6,0 for the non-matching cases)

Mixed strategies: best response correspondences

Stag hunt

	Stag	Hare
Stag	10,10	0,5
Hare	5,0	5,5

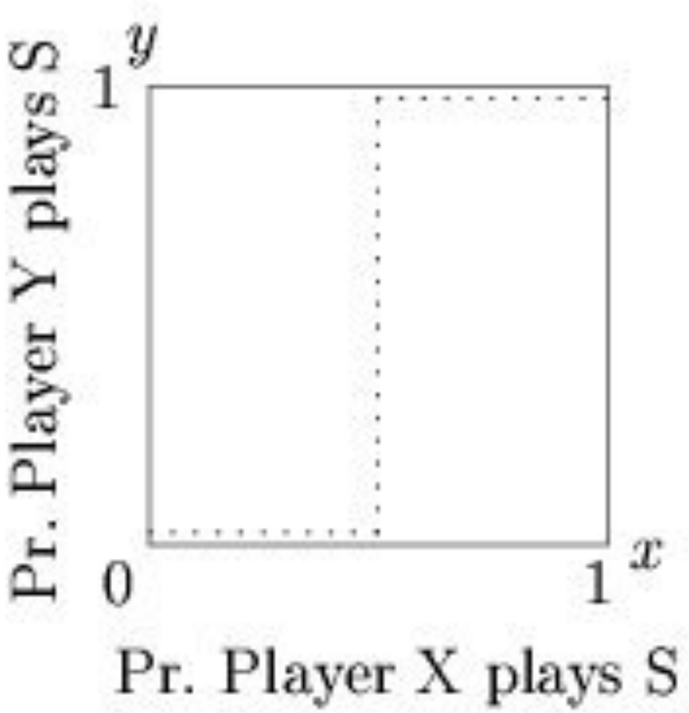

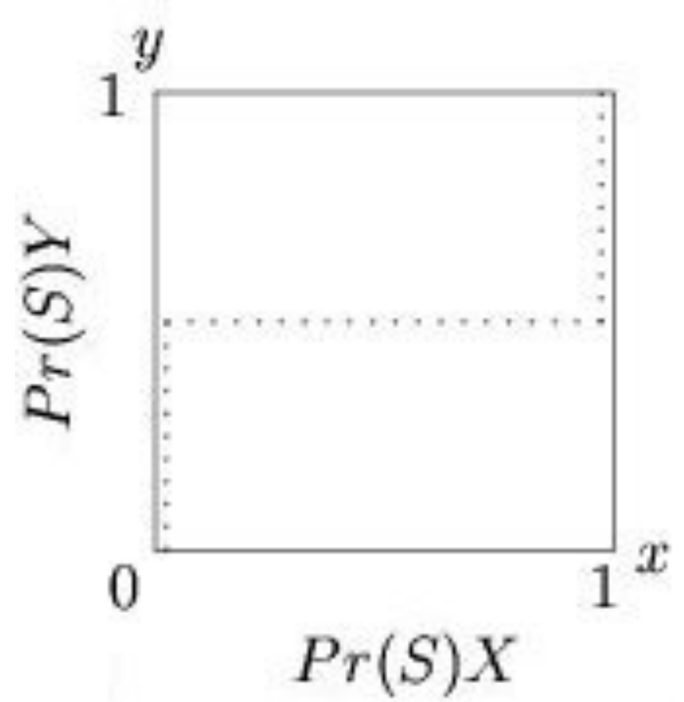

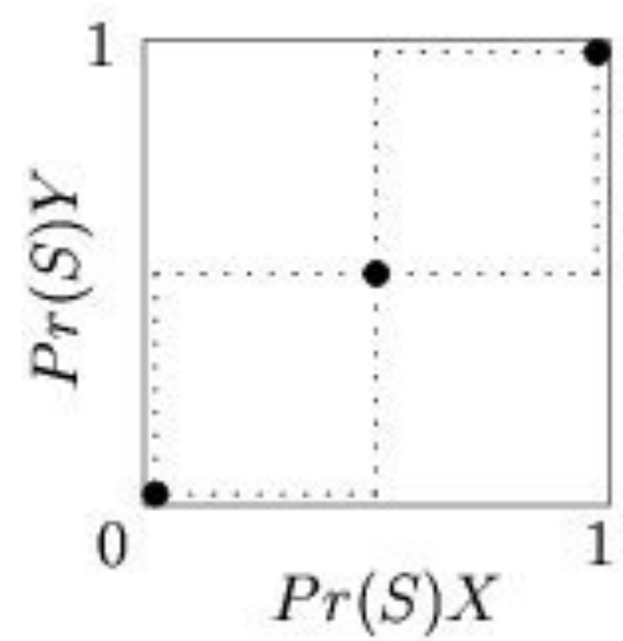

Matching pennies game

		Odd Heads	**Odd** Tails
Even	Heads	+1,-1	-1,+1
Even	Tails	-1,+1	+1,-1

Zero-sum game
Nash equilibria?

Matching pennies game

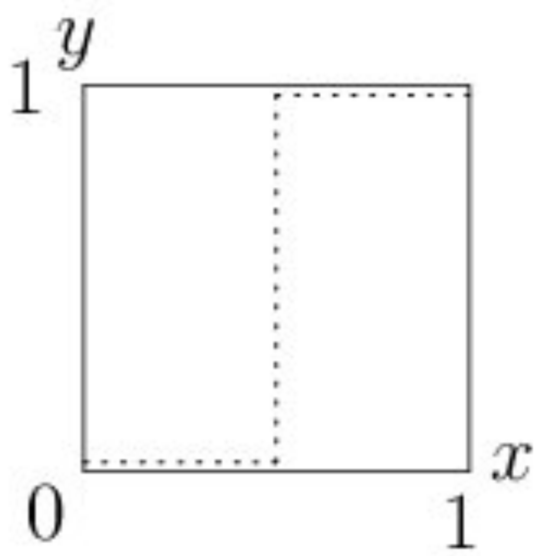

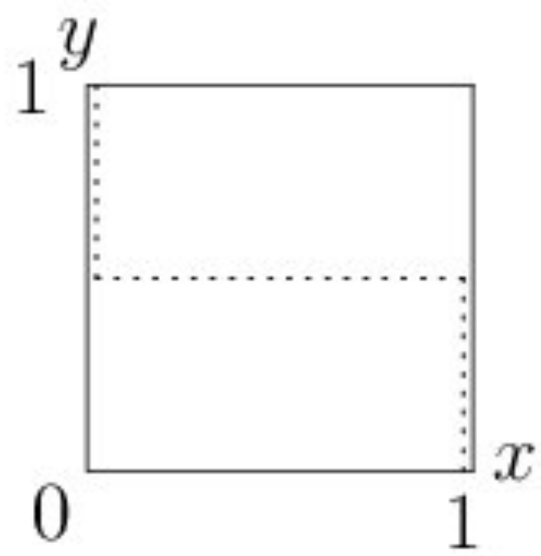

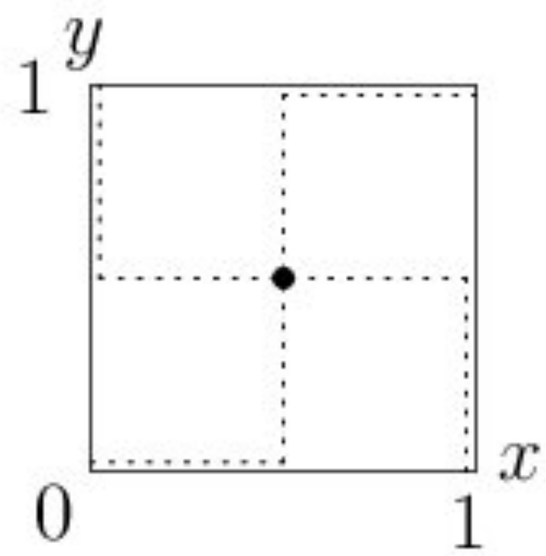

Challenges of games

- The need to remain analytically tractable makes it difficult to incorporate certain aspects of real world circumstances into games:
  - Temporal evolution of populations
  - Stochasticity
  - Space and interaction topology
  - Explorations of heterogeneity
  - Multiple or no equilibria situations

Challenges of games

- There have been many successful analytical approaches to tackling some of these issues (e.g. evolutionary game theory, etc.)

- Given concerns with things like **heterogeneity, space, interaction topology, simplistic actors, adaptation,** and **temporal dynamics**, seems like computational modeling may be useful in understanding these dynamics as well…

Game theory & ABMs

- Basic mapping:

Player Level

Players →	***Agents***
Strategies and Actions →	***Agent rules***
Player Payoff →	***Agent variable***

Game Level

Payoff Structure →	***Global variable called during agent interaction***
Iterations →	***Time Steps***
Number of Players →	***Number of agents involved in an interaction***

Game theory & ABMs

- Capture bounded rationality with agents using simple behavioral rulesets based only on local information
  - **Bounded rationality** - individuals have limited information, cognitive limitations, and finite time to make a decision
- Capture rudimentary “learning” through incorporation of agent memory in behavioral rules
- Introduce interaction topologies to determine who plays (interacts) with whom

Population dynamics & game theory

- Link agent payoffs to fitness and begin with a heterogeneous mix of agents imbued with different strategies

- Can use a tournament (i.e. multiple rounds of interaction) to assess robustness of different strategies or go further and link payoffs to reproduction in next rounds

- Investigate strategy evolution through allowing strategy "mutations" during reproduction (genetic algorithms, evolutionary game theory)

Population dynamics & evolutionary game theory

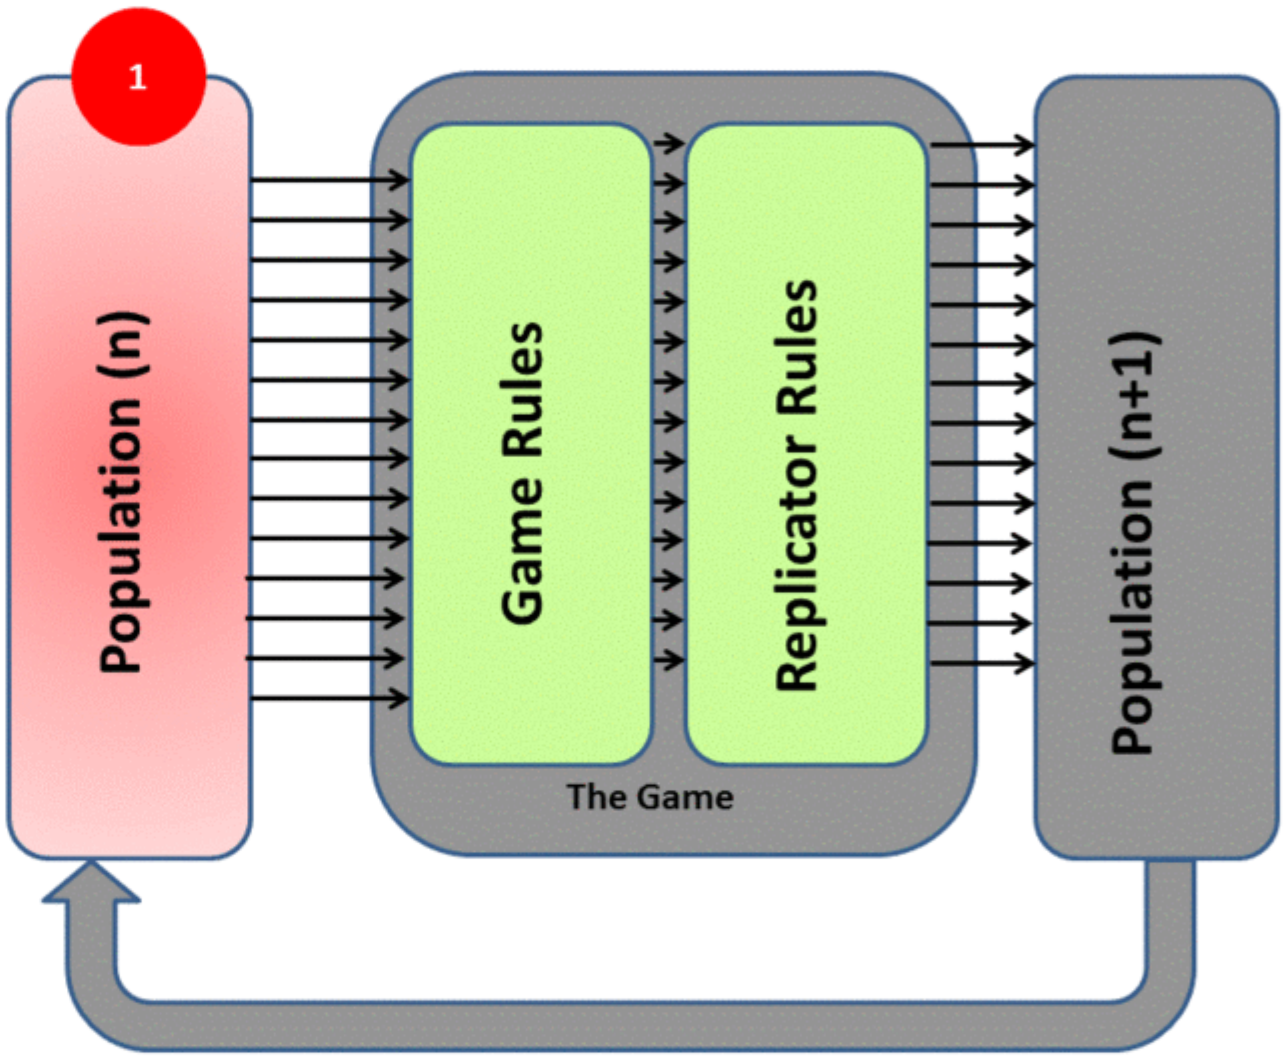

https://en.wikipedia.org/wiki/Evolutionary_game_theory

Evolutionary game theory

- An evolutionary game describes interactions at a single point in time.
- Evolutionary dynamics describe how traits change over time.
- Replicator equation: Every generation, suppose traits increase in prevalence proportional to the difference between their fitness and the average fitness in the population.

Population dynamics & evolutionary game theory

- Continuous time version:

$$\frac{df_i}{dt} = f_i[\phi_i(f) - \bar{\phi}(f)]$$

- where $\bar{\phi}(f)$ is the average fitness

- Wide range of approaches to looking at these issues (replicator-mutator, imitation dynamics, etc.)

- But we can also look at this with agents!

Population dynamics & games

- Note that these don't necessarily have to refer to evolution in a biological sense—evolutionary game theory and similar approaches are often used to understand many different systems, e.g.:
    - Infectious diseases (e.g. disease/behavior feedback loops—consider social distancing, vaccination, etc.)
    - Voting patterns, communication
    - And many other systems where behavior may change over time

Evolution of Cooperation

- Perhaps the most famous example of incorporating game theoretic model into an ABM context comes in classic study of the Evolution of Cooperation (Axelrod and Hamilton, 1981; Axelrod 1984)

- Begins with a persistent problem in both the social sciences and biology:

    - **How can cooperative behavior in groups arise and persist?**

Cooperative behavior and altruism

- Humans, many different animals (bats, etc.)
- However, cheating would often seem to gain higher payoff, why does cooperation and altruism persist?
- Historically, there has been a major debate on how individually costly behavior that benefits the group can arise and sustain within populations.
- Let’s look at this with prisoner’s dilemma

Cooperation and the Prisoner's Dilemma

- Being in a group of cooperators is good, but being a defector in a group of cooperators is even better.

- Holds true for biology as well: if payoffs are linked to reproduction, who will produce the most offspring?

		Prisoner A	
		Silent	Betray
Prisoner B	Silent	-1 -1	0 -3
	Betray	-3 0	-2 -2

Evolution of cooperation

- Axelrod’s Insight:
- In a one-shot PD game, “Always Defect” [All-D] always wins at both the individual and population levels (anything else can always be “invaded” by a newcomers playing of [All-D])
- In an iterated PD game with an uncertain time horizon and a basic ability to remember prior interactions, however, **other strategies may also be potentially stable** (although note that now replicator style approaches are less convenient)

Evolution of cooperation

- Are there simple strategies relying on simple memory that can allow cooperative group behavior to succeed in situations of on-going interaction?
- Success Criteria
  - Robustness: thrive in mixed population of strategies
  - Stability: once established can resist “invasion”
  - Initial viability: can establish in the midst of a lot of Defectors

Axelrod’s tournament

- **Agents**
  - Agents are assigned to play one of 14 extremely simple to somewhat more elaborate strategies submitted by a set of experts
  - Strategies also include [All D], [All C], and [Random]

Axelrod’s tournament

- **Model Setup**
    - Round-robin tournament of one-to-one matchups of all strategy pairs
    - Each matchup goes for 200 iterations (but agents don’t know that)
- **Model Outcome Assessment**: see which strategy had the highest average payoff across whole tournament

Axelrod’s tournament

- The Winner:
  - **Tit-for-Tat [TFT]**
- Even though extremely simple and involving only a very short memory, [TFT], that involves basic “nice” reciprocal cooperation, won out over everything else – including [ALL D]!

Let’s play!

- Fantastic version of this by Nicky Case
- https://ncase.me/trust/
- Netlogo iterated prisoner’s dilemma

Axelrod's tournament (Round 2)

- **Agents**
  - 64 more strategies submitted from experts in a large number of fields (including Game Theory)
- **Model Setup**
  - Same round-robin tournament of one-to-one matchups of all strategy pairs
- Also looked at an "ecological" variant where populations for the next tournament were proportional to success in prior tournament (generated a time path of strategy distributions)

Axelrod's tournament (Round 2)

- The Winner:
    - **Tit-for-Tat [TFT]** (Again)
- Here too, this basic strategy dominated both in terms of average success AND by completely taking over the population distribution in the "ecological variant"

However—

- Tit-for-tat success can be very sensitive to exactly how the tournament/game is structured
- Possibility for a “death spiral” - e.g. if occasional random errors in choosing
    - Tit-for-two-tat, etc.

Let’s try a few examples!

- https://colab.research.google.com/drive/139U20G4NHs-DSOeZeIUVZHthIDSeTz39

Take-home messages

- Given a set of extremely plausible assumptions (like some initial clustering of cooperatively inclined individuals in a population), the basic principle of reciprocal cooperation can outperform an "All Defection" approach

- Without any appeals to "group selection," can explain from "the bottom-up" emergence and persistence of cooperative behavior

- Given importance of bounded rationality, heterogeneity, and temporal evolution of populations in this analysis, very unlikely we could have gotten these results without availability of computational modeling

Another example: the Diamond & Dybvig model of banking

This example is mostly borrowed from John Leahy's excellent talk on the Economics Nobel prize from the Complex Systems 2022 Nobel Symposium!

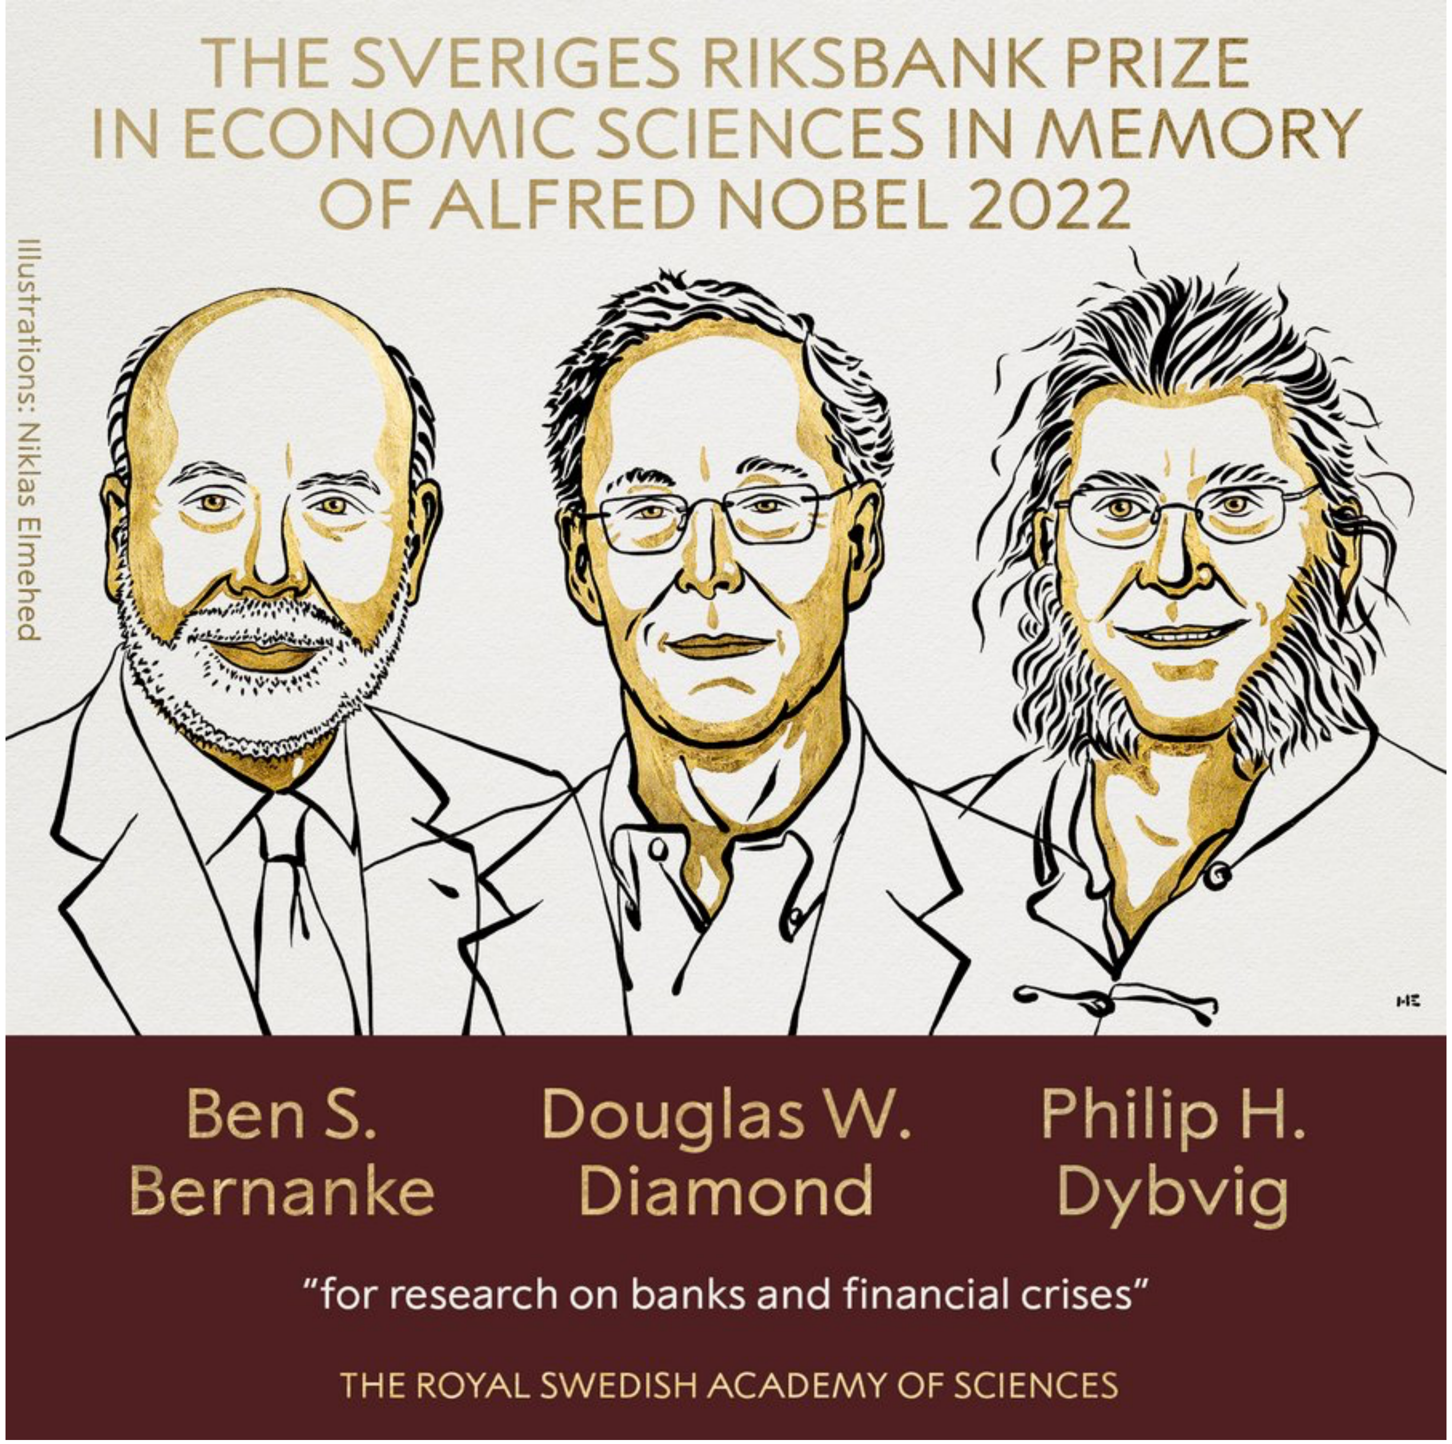
THE SVERIGES RIKSBANK PRIZE
IN ECONOMIC SCIENCES IN MEMORY
OF ALFRED NOBEL 2022
Illustrations: Niklas Elmehed
Ben S.
Bernanke
Douglas W.
Diamond
Philip H.
Dybvig
"for research on banks and financial crises"
THE ROYAL SWEDISH ACADEMY OF SCIENCES

Banks play an important role in our economy—they provide means for payment and financing

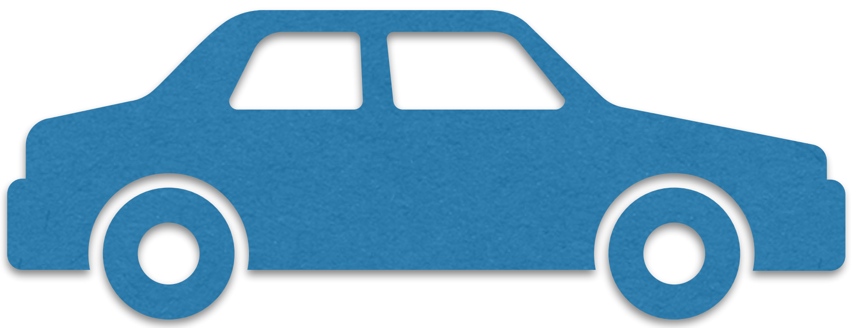

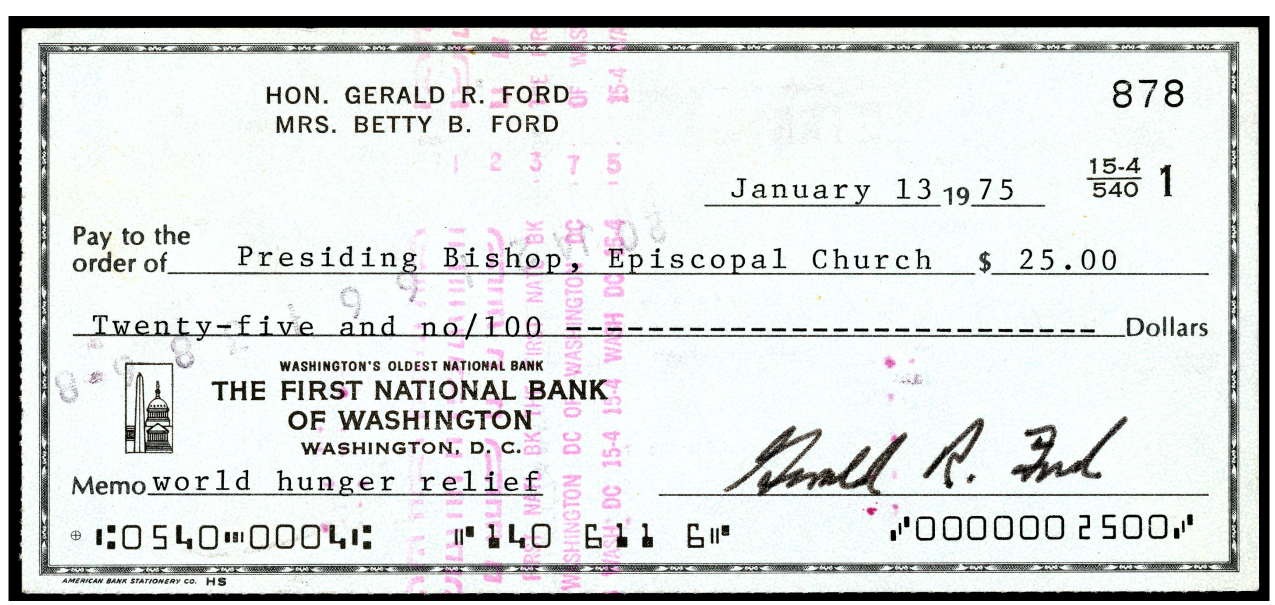

But their vulnerabilities also lead to economic crises

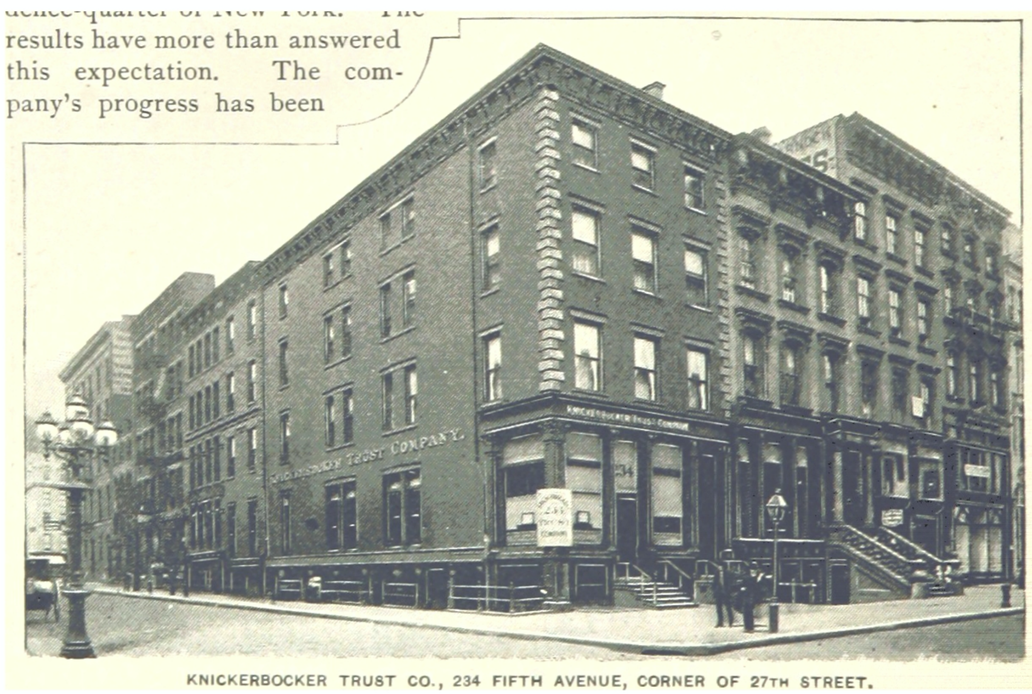

FTX

How do banks work?
A very simplified view

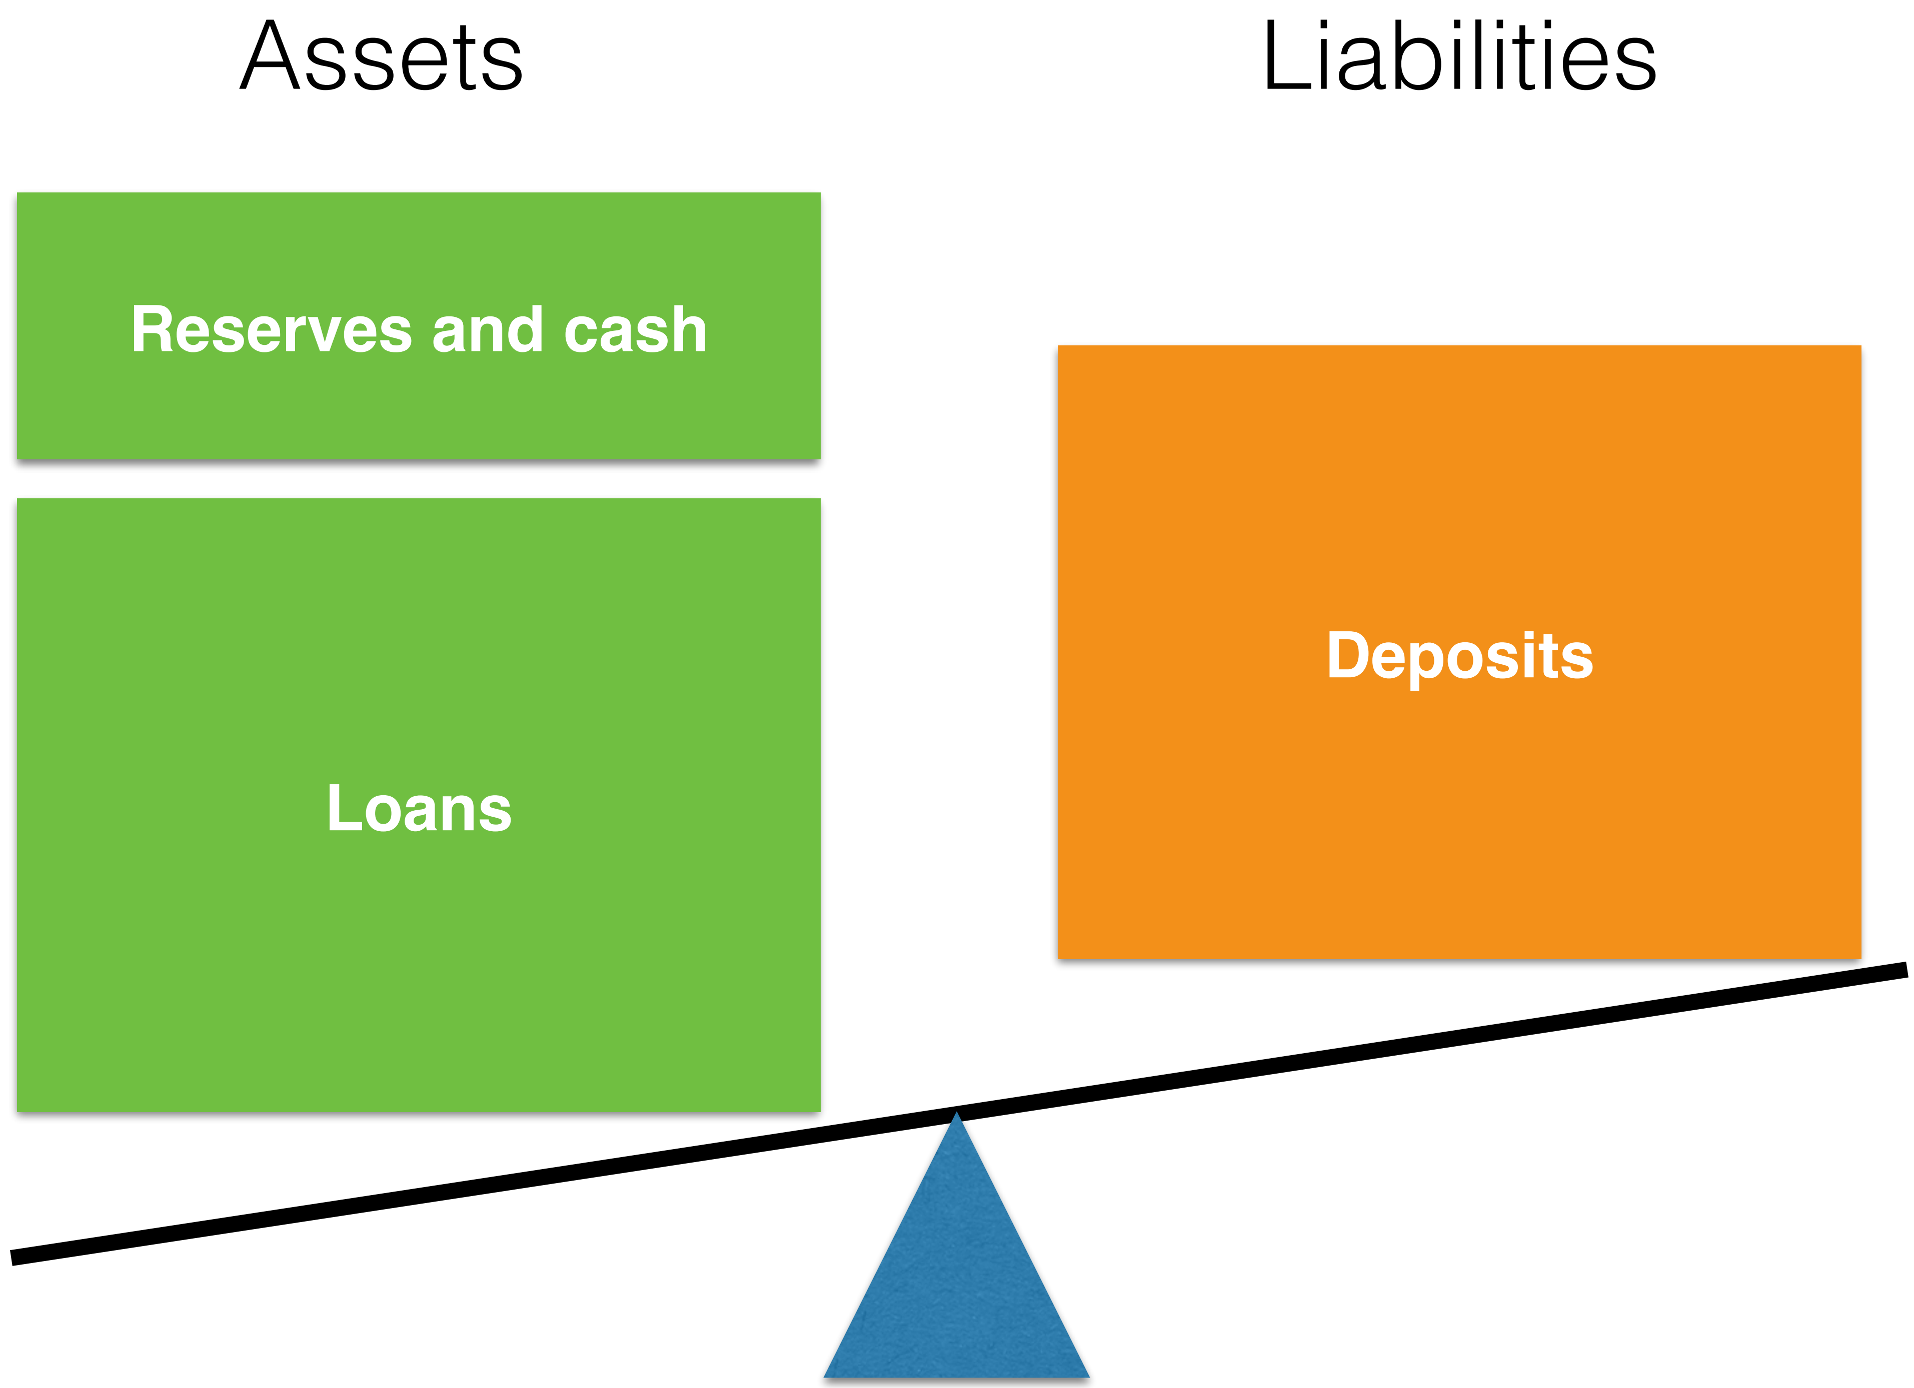

But, deposits are liquid and loans are not—banks bet that most people won't need liquid assets at the same time

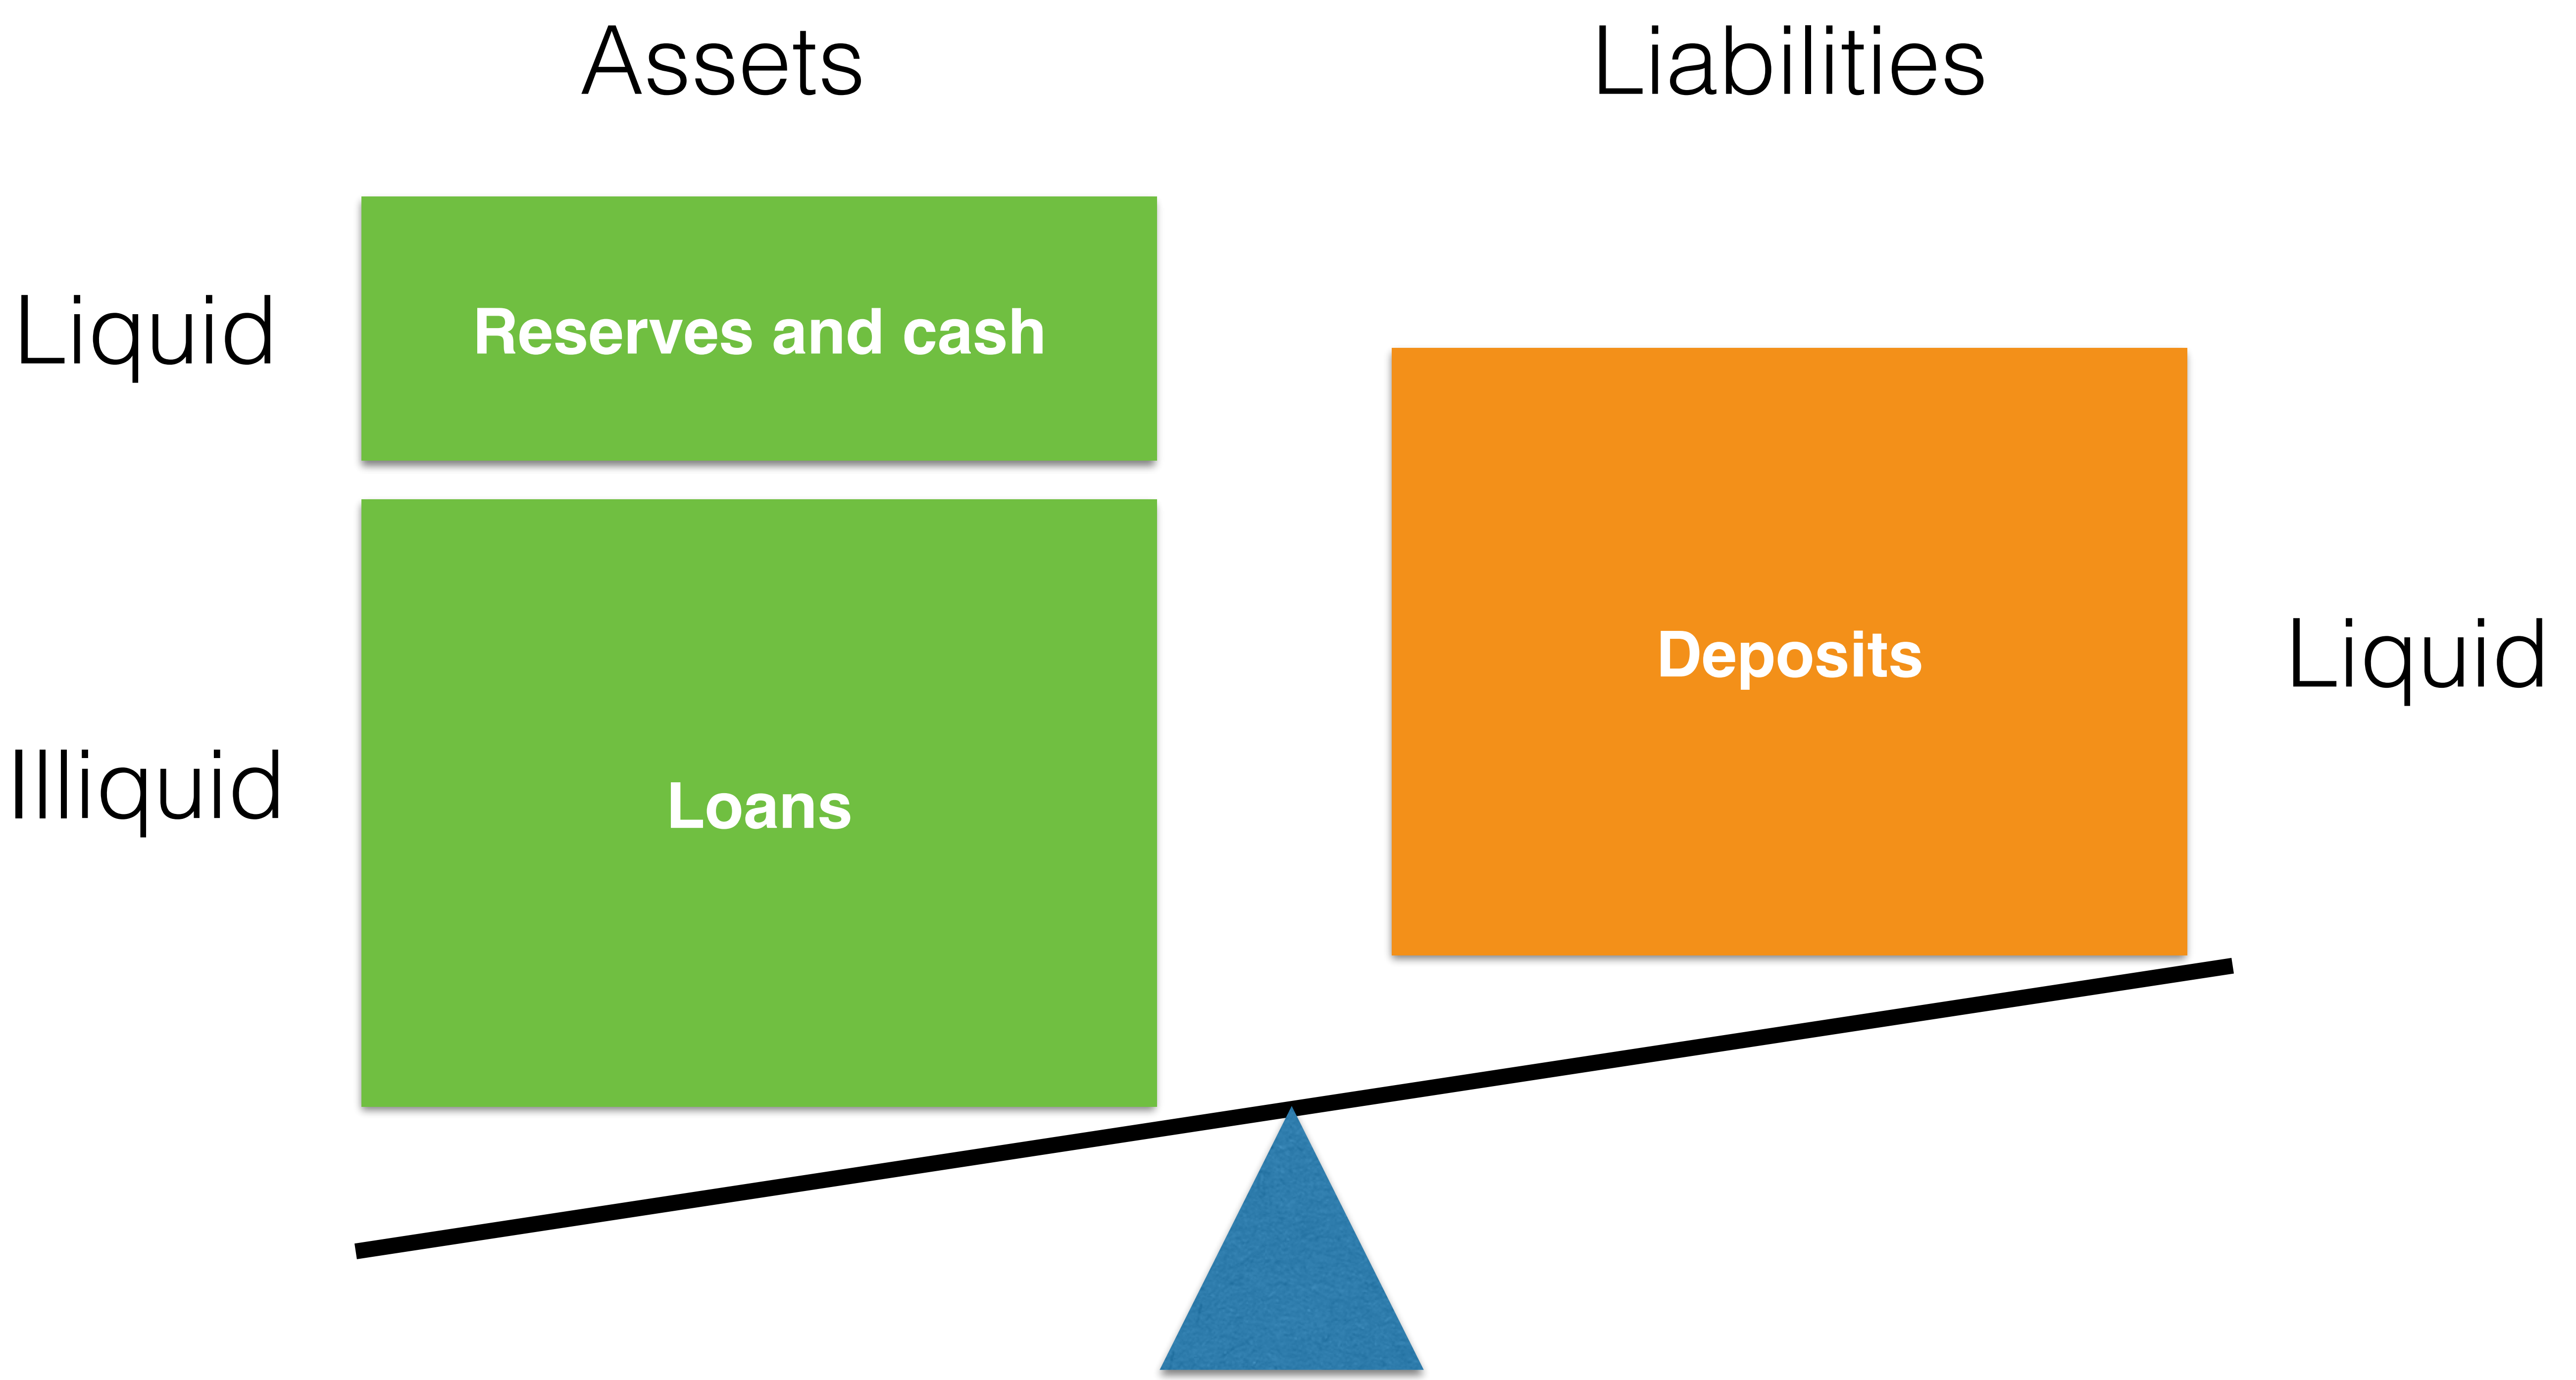

Insolvency: if the value of the loans is reduced then concerns that assets may be < deposits—leads to fear in depositors

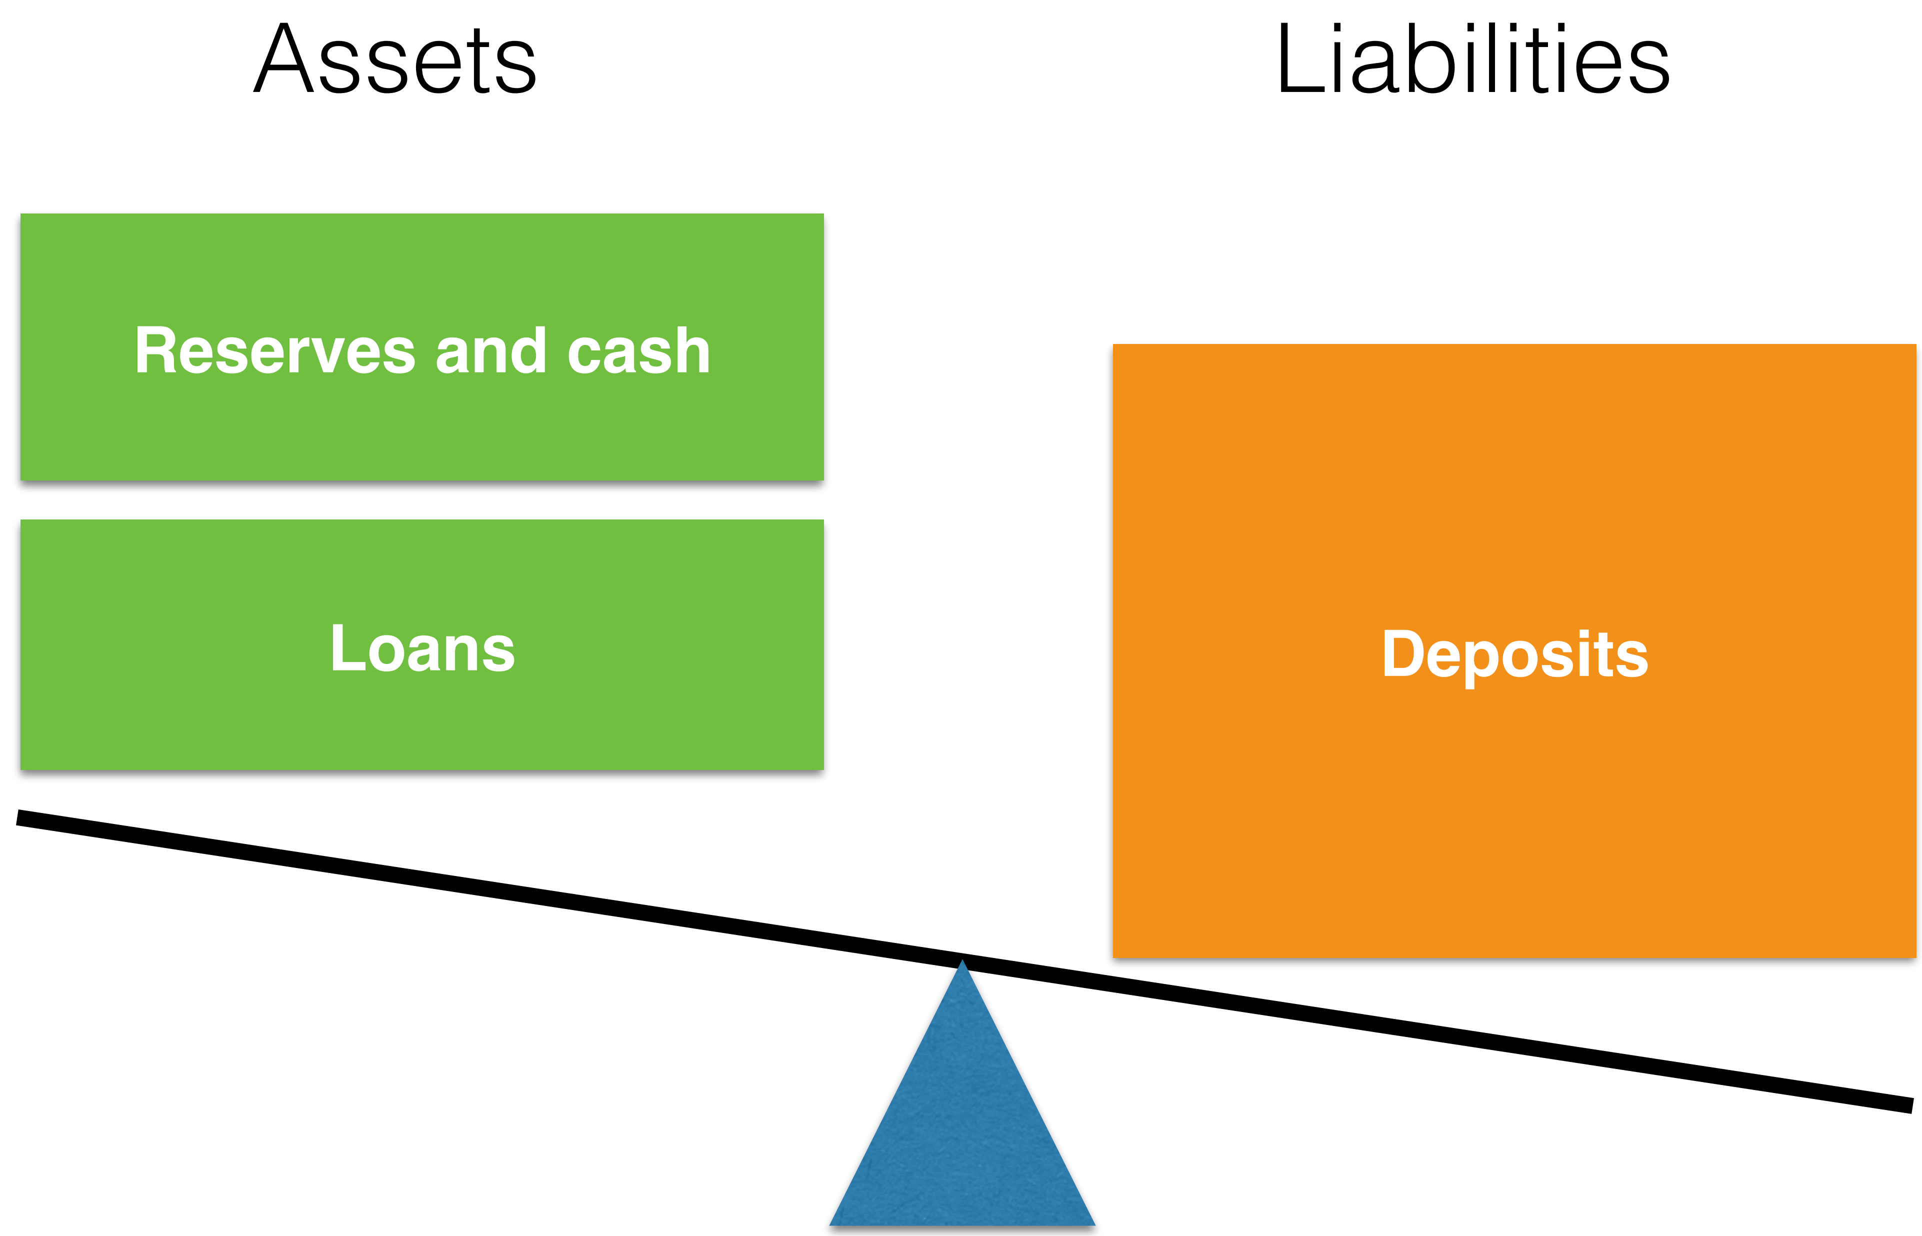

The result?
A bank run

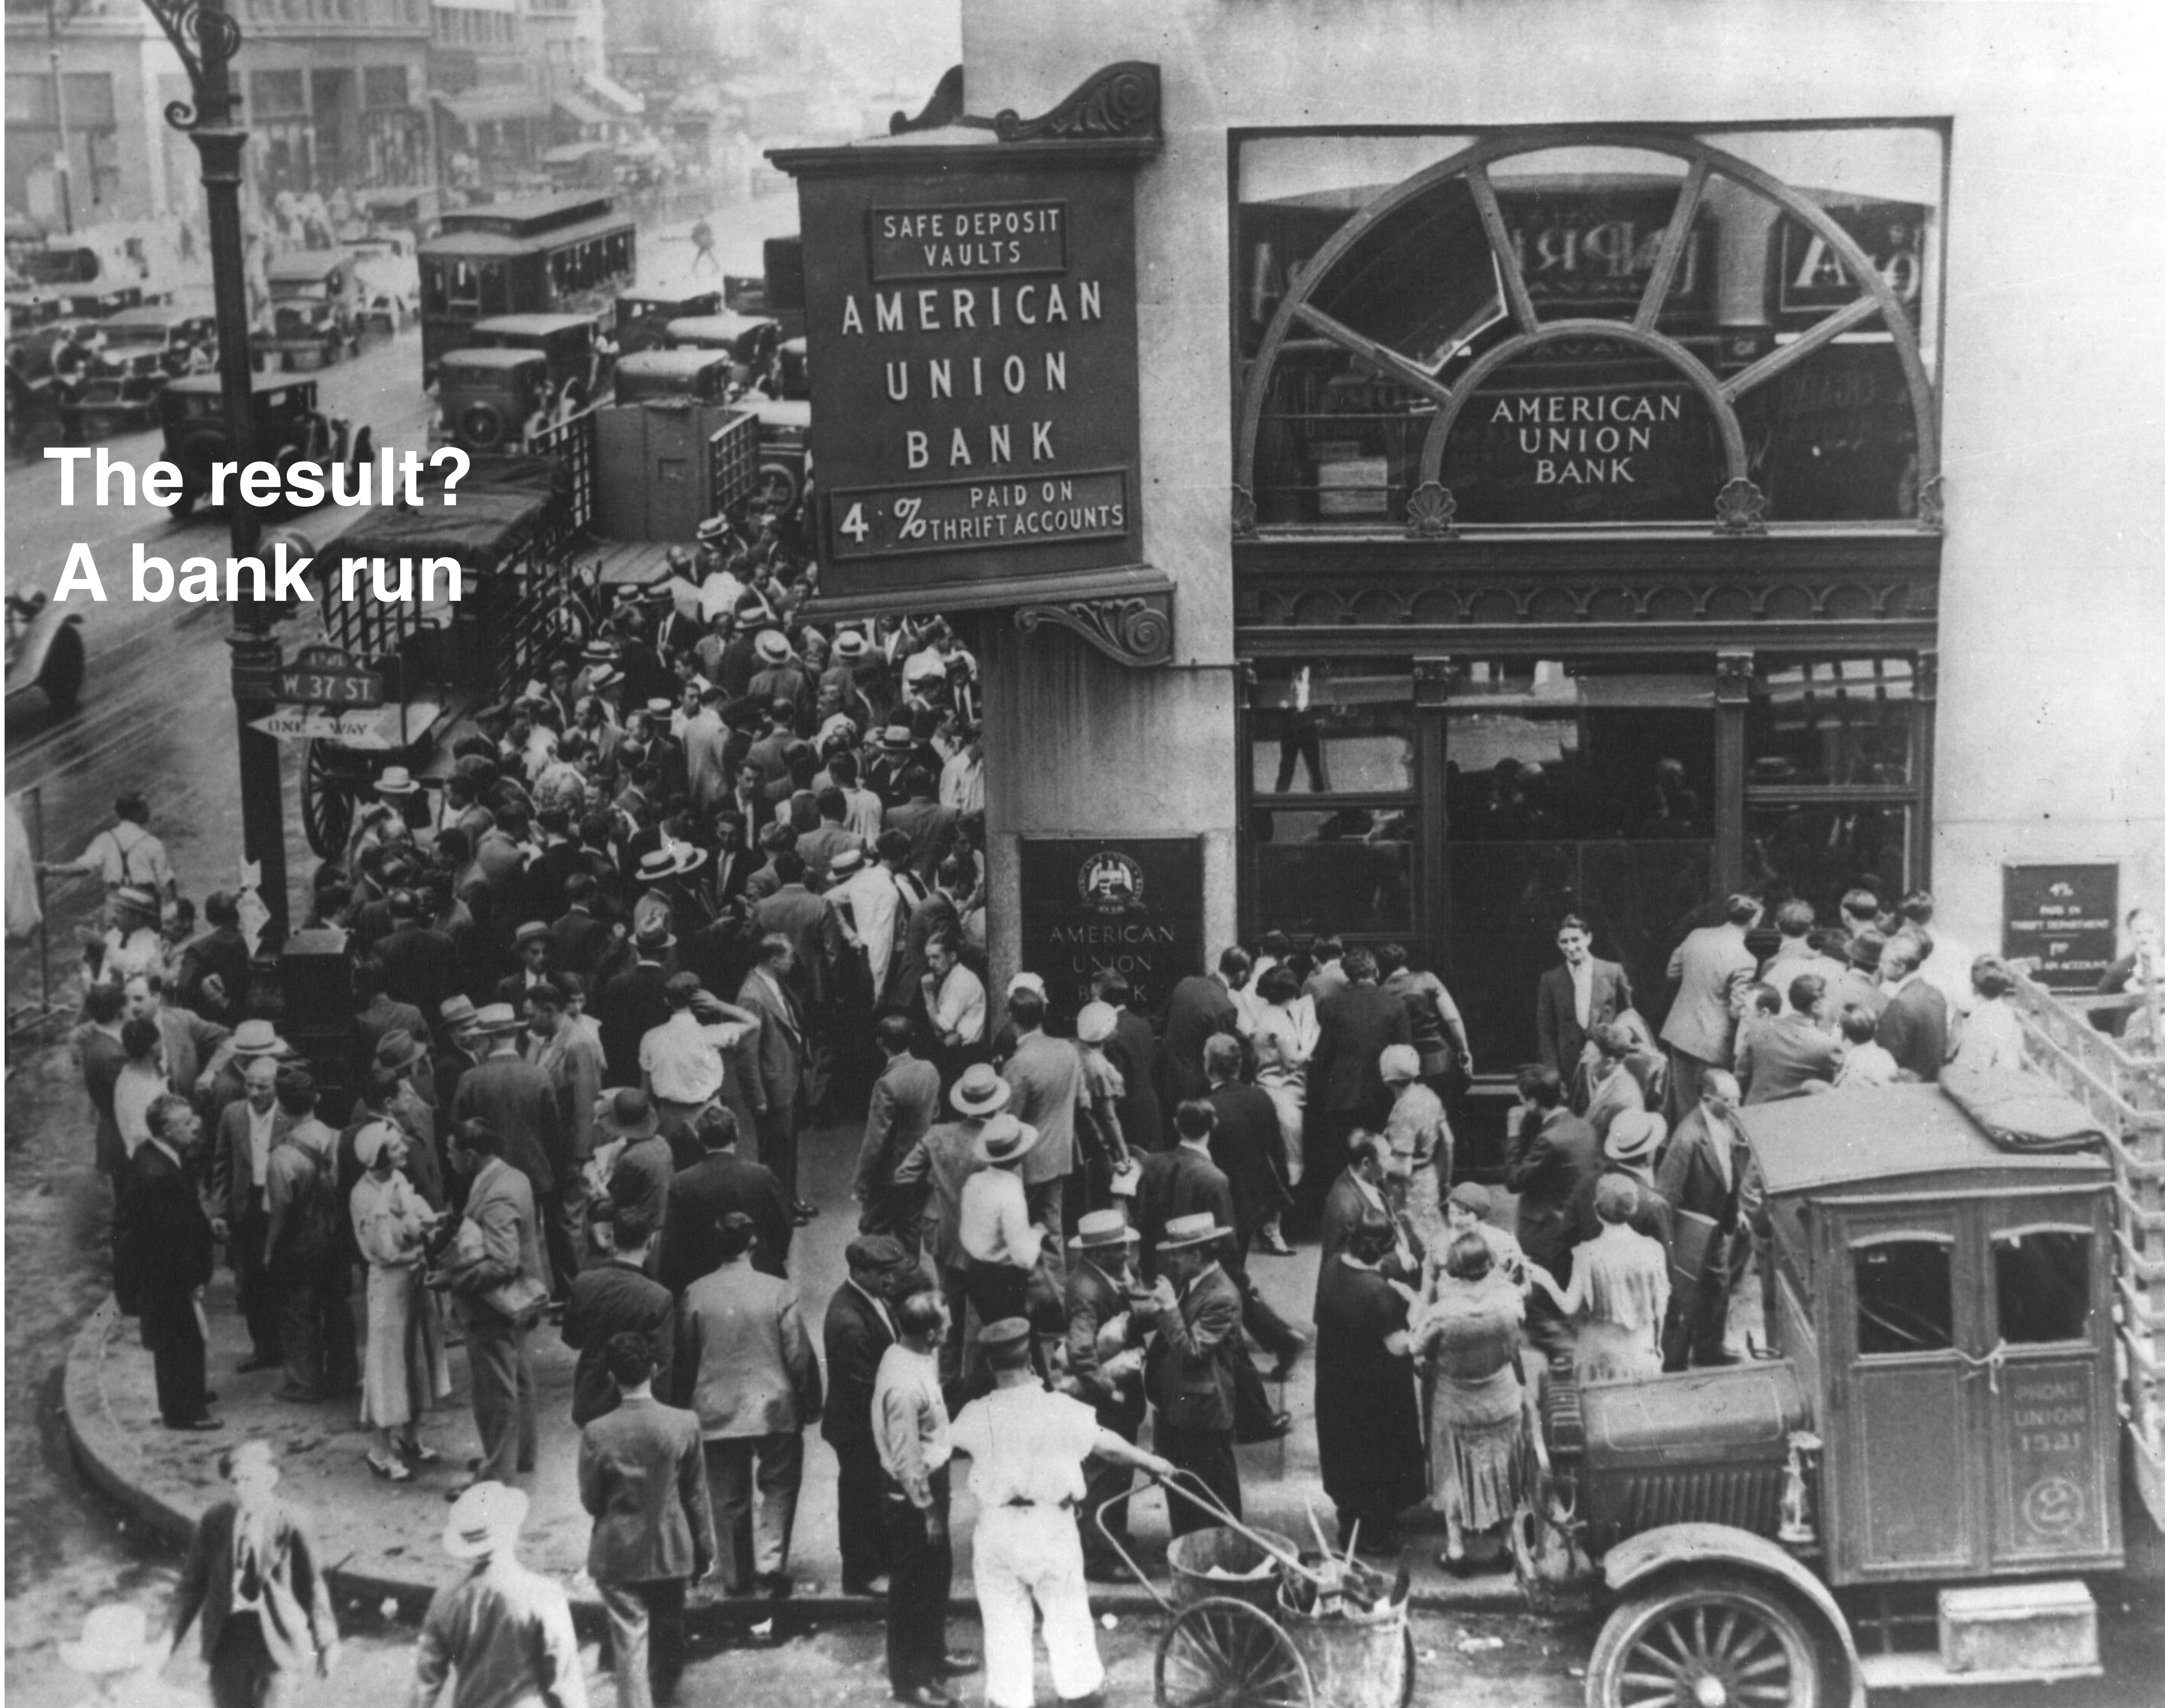

Diamond & Dybvig model

- Many, many variations, we'll do a simple one here
- Explores questions of: why do we have banks? Why is liquidity needed? What is the role of expectations and how does financial contagion work?

Diamond & Dybvig model

- Time: there are 3 time periods, t_1, t_2, and t_3
- Agent properties:
  - Each agent starts with 1 unit of resources that they wish to save for future consumption need
  - Each agent has a consumption date where they will need consume 1 resource.

Diamond & Dybvig model

- Agent Properties
  - Some will need to consume their resource at t_2 and others at t_3—but the agents don't initially know their type! (to represent that we can't anticipate when there will be an emergency and we need resources)
  - Suppose a 50% chance of needing to consume at t_2 vs t_3
  - Agents learn their type at time t_2

Diamond & Dybvig model

- Agent actions and rules:
  - **Storage**: the agent can just store the 1 resource for use later (no change to resource level)
  - **Long term investment**: the agent can invest the resource at time t_1 and receive $L<1$ if at t_2 or $H>1$ at t_3
    - If they have to liquidate early there is a cost and they get less than they started with (L)
    - But, if they can stick with the investment until t_3 they gain more resources (H)

Diamond & Dybvig model

- Storage is liquid but does not pay interest (could even consider a variation where it depreciates)
- Investment is illiquid but pays a higher return if held to completion

What if no banks?

- Agents want to maximize their expected resources
- The payoff to storage is $1/2 \cdot 1 + 1/2 \cdot 1 = 1$
- The payoff to investment is $1/2 \cdot L + 1/2 \cdot H$
- Which to choose depends on L, H, risk aversion
- Ideally want to choose storage if you are going to consume at t_2 and investment if you are going to consume at t_3—but no way to know which
- Banks make this possible

Banks

- Suppose non-profit banks with no costs and give back all returns to depositors
- Banks offer deposit contracts
  - Pay preannounced rate of interest based on investment project
  - Depositors can receive the full value of their deposits at any time **so long as the bank has funds**

Banks

- Banks return resources/funds to depositors in the order they request them for as long as they have funds
- Suppose we randomize for all who request at the same time point, to simulate random arrival times at the bank
- The bank can take advantage of scale—individuals cannot know their consumption type but the bank knows that overall it will be 50% of each and can plan/invest accordingly

So what happens?

- This relatively simple model has many different Nash equilibria!

The good equilibrium

- Everyone deposits their money in the bank
- The bank keeps half of deposits in storage and invests the other half
- All of the early consumers withdraw at t_2 and receive one unit of resource
- All of the late consumers withdraw at t_3 and receive H units of resource
- Everyone gets to do what they wished they could do in the no-banks scenario!

But alternatively…

- If everyone else is going to withdraw, the better strategy is to withdraw, to avoid the bank running out of funds before you need to consume (e.g. if they have to liquidate investments and only get L instead of 1 resource)
- So another Nash equilibrium can be that everyone withdraws—a self fulfilling prophecy of sorts
- For example suppose—

Bank run equilibrium

- Suppose that there is a public signal that all can see (that has nothing to do with the banking system—a “sunspot”)
- Suppose that depositors believe that other depositors will withdraw their funds if the signal is observed
- Suppose that this signal occurs with a very low probability.
- Then initially everyone acts as in the good equilibrium
- At t_2 if no signal then everyone acts as in the good equilibrium, but if the signal is observed all agents attempt to withdraw their funds
- And we have a bank run!

Diamond & Dybvig

- In this model, banks arise because they provide higher returns and more liquidity than an individual can get on their own.
- The mismatch of between liquid liabilities and illiquid assets gives rise to multiple equilibria and instability

Belief/market sentiment is important: beliefs can become self fulfilling

- A bank run can occur even when banks are sound!
- President Roosevelt 1933: “The only thing that we have to fear is fear itself.”—he was talking about bank runs
- Knickerbocker trust bank run in 1907—everyone eventually got their money back!
- Stop bank runs by restoring confidence

How to fix it?

- Deposit insurance can largely eliminate bank runs by removing the need to “get there first”
- Since the FDIC in 1930’s, we haven’t really seen runs on the commercial banking system (talk about Silicon Valley Bank)

The Great Recession

- Often the liquidity of assets changes, especially during times of crises
- During the Great Recession, banks held lots of mortgage backed securities (MBS)
- These assets had been very liquid prior to the crisis, but the market dried up as mortgage defaults rose and their quality came into question—they became illiquid

The Great Recession

- Essentially a run of the banking system on itself (on the "shadow banking system")—a big bank run of banks on each other
- Failure of Lehman created uncertainty about other financial institutions (AIG)
- The Fed responded by providing liquidity—lending treasury bills against bank collateral at high interest rates… the so-called "bailout"
- The Fed made a $52 billion profit in 2009

Financial contagion

- Why did bank X go under? Is my bank vulnerable?
- Failure of one bank can drive down asset prices, causing other banks to have questionable net worth
- FTX when they sold off crypto—fire sales

For next time…

- The evolution of trust: https://ncase.me/trust/
- The Evolution of Cooperation, Robert Axelrod; William D. Hamilton. Science, New Series, Vol. 211, No. 4489. (Mar. 27, 1981), pp. 1390-1396.