Artificial intelligence and systemic risk

Jon Danielsson
Robert Macrae
Andreas Uthemann

Systemic Risk Centre

modelsandrisk.org/AI

16 May 2019

From

- modelsandrisk.org/AI
- Artificial intelligence and the stability of markets
- SRC discussion paper
- voxeu.org/article/artificial-intelligence-and-stability-markets

Artificial intelligence (AI)

- Take the
  - Machine learning (ML) associations
  - rulebook
  - supervisor interface with the regulated institutions
- Have the AI identify how to best achieve supervisory objectives
- Suggest or make supervisory decisions

What AI can and cannot do

- AI can master any decision process with a *defined action space* better than any human
  - chess, go, , computer games,...
- If the action space is ill defined (like all human endeavours)
- AI today is *unable to reason about things it has not seen*
- It can generalise within a local problem but cannot apply experiences from one domain to another
- Because it does not understand the global problem in which the local one is embedded
- It can handle decisions to the extent they can be mapped onto a *contained local problem*
  - driving a car, medical diagnosis, allocation of credit

Bob, the Bank of England Bot, and friends

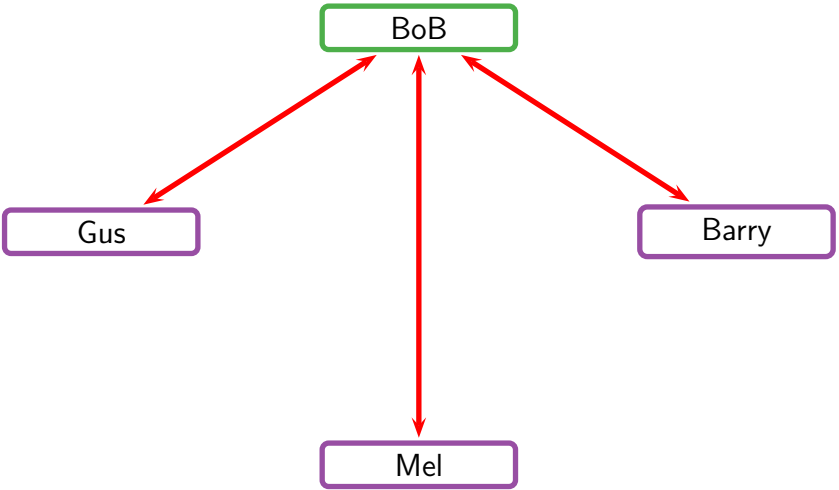

Bob, the Bank of England Bot, and friends

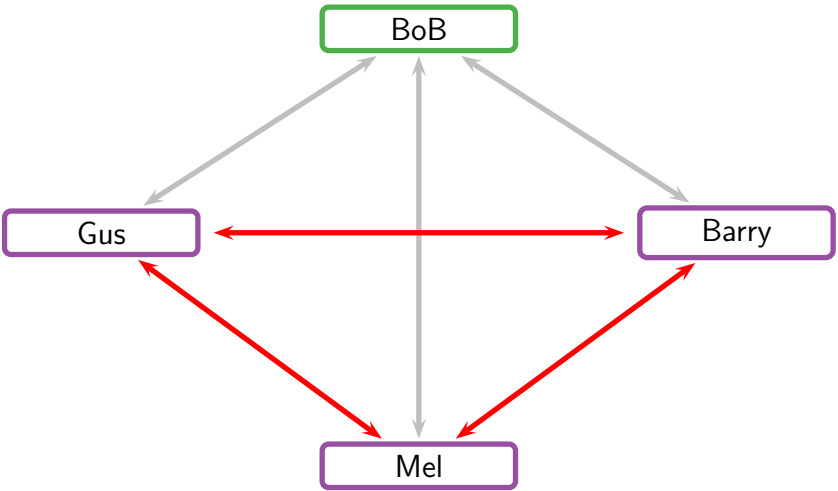

Risk management, compliance and micropru

- Prime candidates for AI
- Most risk modeling as currently done can be outsourced to AI
- Just like much of the rest of risk management and micropru
- Very significant cost and efficiency savings
- Opposition is social, political, legal but not technical
- Project Mason
- FCA rulebook is now machine readable logic engine with a bot interface

The time dimension of risk

The time dimension of risk

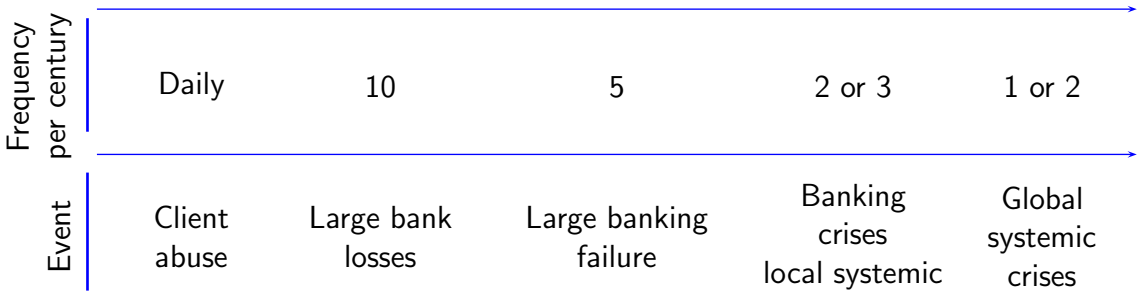

The time dimension of risk

Frequency per century	Daily	10	5	2 or 3	1 or 2
Event	Client abuse	Large bank losses	Large banking failure	Banking crises local systemic	Global systemic crises
Drivers	Profits	Idiosyncratic risk	Systemic risk	Macro economy	Politics

The time dimension of risk

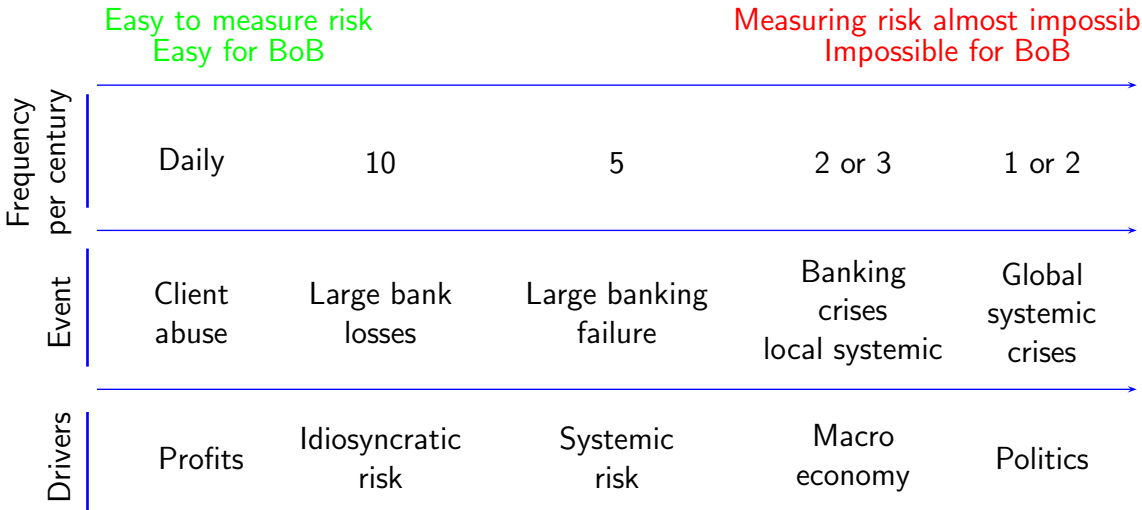

What can go wrong?

1. AI can't reason about things it has not seen
2. And is unable to deal with unknown–unknowns
3. While it is procyclical
4. And easy to attack

Inability to do causality and reason

- A 1980s AI, EURISKO, played a naval wargame
- It found the best solution was to sink its own slowest ships
- It is impossible to specify all eventualities
- Humans can reason about unseen things, AI will not
- But AI will make decisions, so it will need a *kill switch to prevent it from doing something catastrophic*

The need for a kill switch

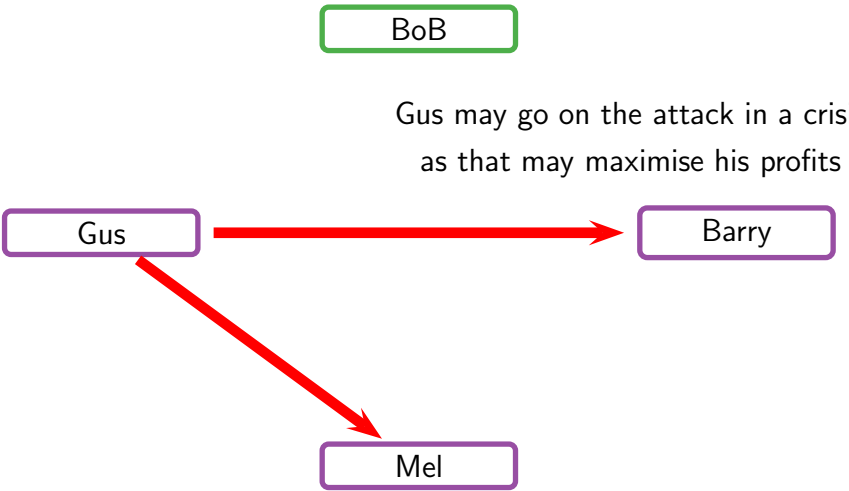

The need for a kill switch

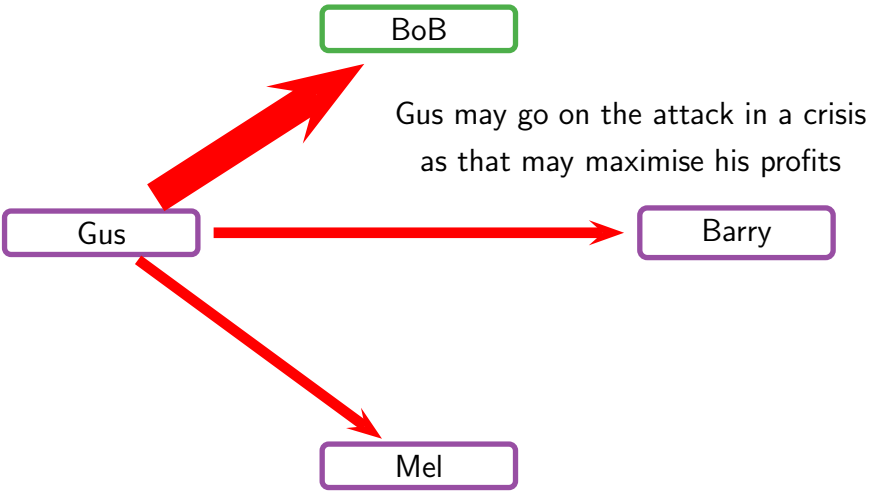

Procyclicality

- AI will favour homogeneous best–of–breed methodologies and standardised processes even stronger than human authorities
- In-breeding and homogeneity will make behaviour more procyclical
- Which increases systemic risk

BoB cannot find unknown–unknowns

- Systemic vulnerabilities tend to happen on the boundaries of areas of responsibilities — silos
- Where we are least likely to look
- In a system that is endogenously infinitely complex
- The machine cannot be trained on events that haven't happened yet
- Therefore, it would be very good at known–unknowns
- And miss the unknown–unknowns that cause crises

Optimise against the system

- It is easier to optimise against BoB than human regulators because
- BoB faces an infinitely complex computational problem
- A hostile actor only has to optimise against very small part of that domain
- Standard responses from AI systems, such as a *randomised responses*, are not acceptable

Macro problems

- To be effective, the macroprudential AI needs to
  1. control across borders
  2. control across silos
  3. share data across borders and silos
  4. randomise responses
  5. create rules in a nontransparent way

  6. understand causality in in unforeseen cases
  7. reason on a global rather than local basis
  8. identify threats that have not yet had bad outcomes
- The first 5 are unacceptable; the last 3 are beyond current capabilities

So...

- BoB and his friends will become increasingly useful to microprudential regulators and risk managers
- Reduce costs for financial institutions and supervisors
- Change the job of the supervisor
- Increase systemic risk
- Reduce volatility and fatten tails

Low vol — Fat tails

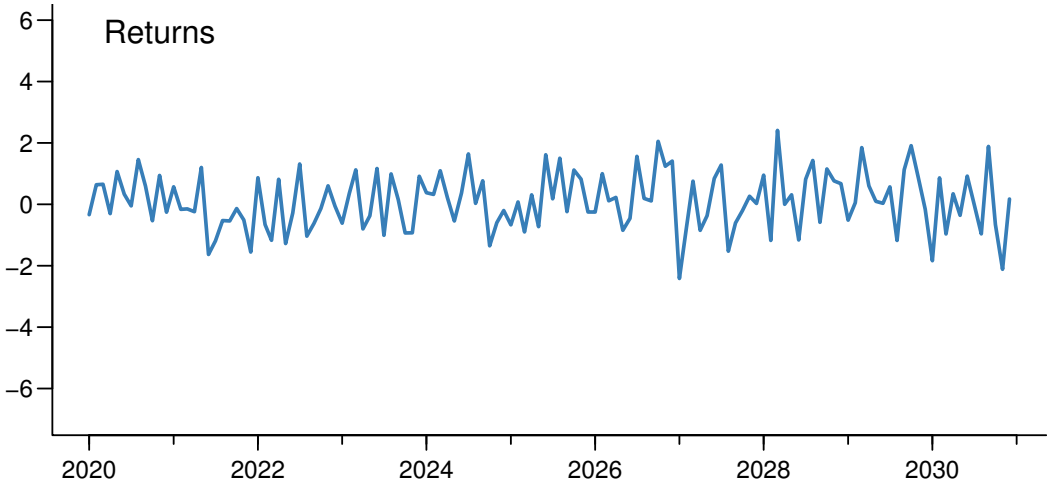

Low vol — Fat tails

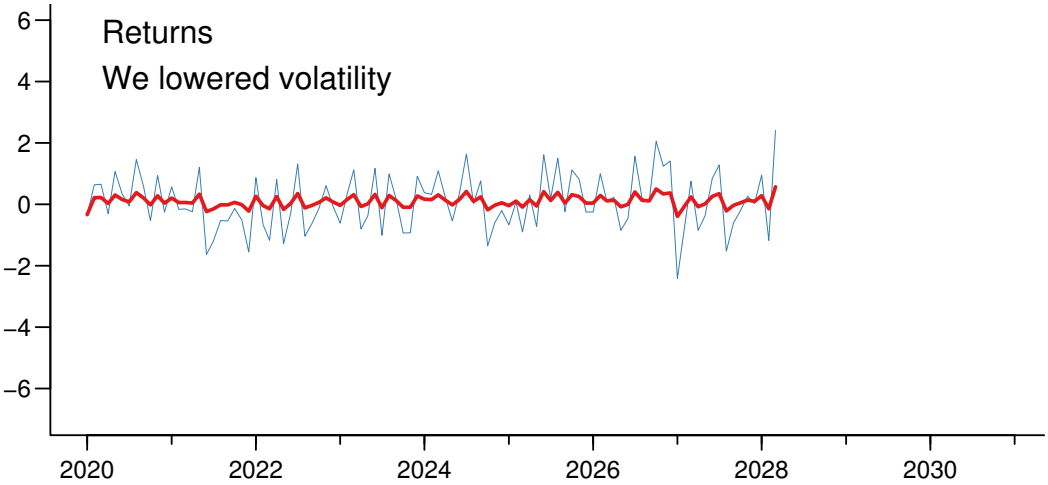

Low vol — Fat tails

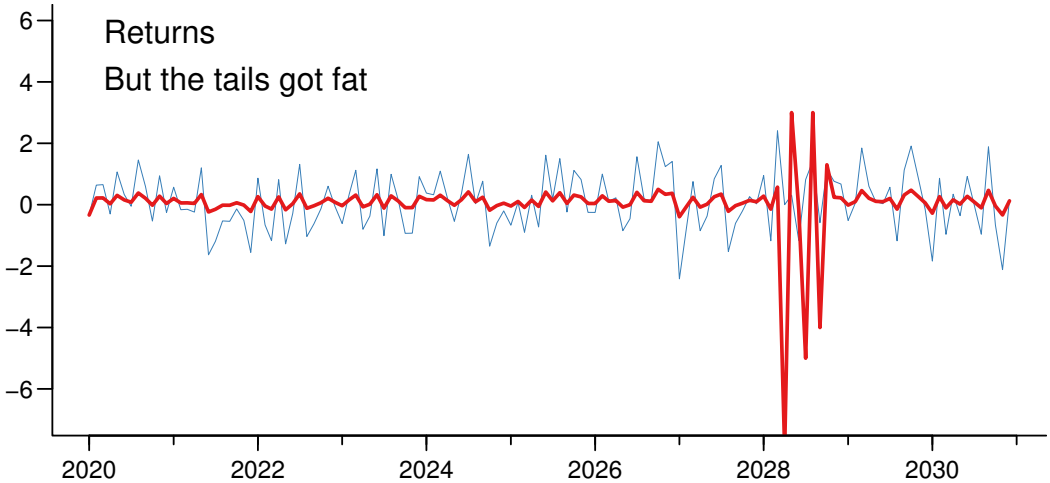

Low vol — Fat tails