Littler

Federal Reserve Bank of Atlanta
Annual Conference on Financial Markets

Public Risk Management for AI:
The Path Forward

May 8, 2018

Matthew U. Scherer, J.D.
Littler, Portland
mscherer@littler.com
503.889.8881

Littler
Employment & Labor Law Solutions Worldwide®

Historical Background: Pre-Industrialization

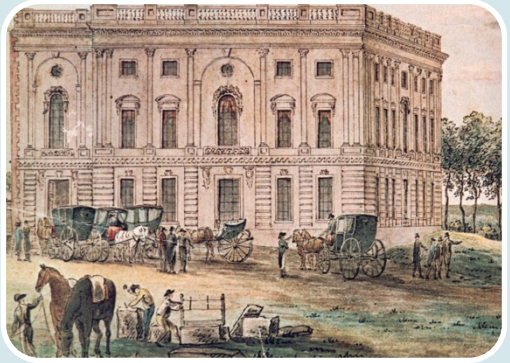

Economy: Agricultural

Decentralized production

Society: Rural

Law: Informal

Legislatures met infrequently

No specialized agencies

No large-scale risks, so risks mostly managed through resolution of individual disputes

Historical Background: Industrial Revolution

Post-Industrial Revolution

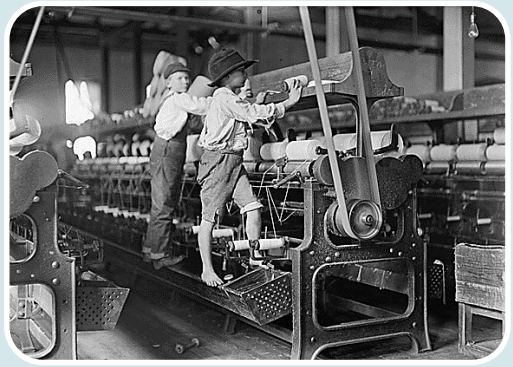

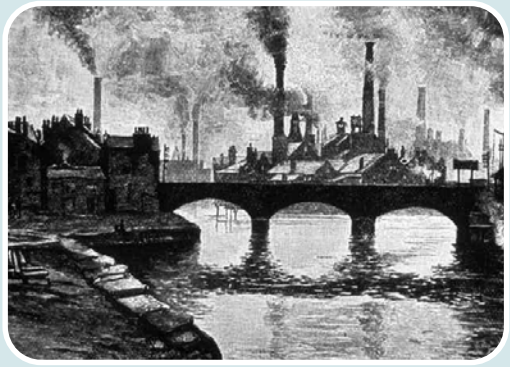

Economy: Industrialization

Centralized production

Society: Urbanization

Law: Formalization

Historical Background: Industrialization's Impact on Law

- **Industrialization created new challenges for the legal system**
  - Defective mass-produced products
  - Workplace hazards
  - Environmental threats
  - Large, powerful private companies that could dominate entire industries
- **Existing legal mechanisms were unable to cope with the effects of these new *public risks***

What is a “public risk?”

A potential source of harm that is:

1. Centrally or mass-produced or widely distributed

and

2. Outside the control of the individual risk bearer

Plain English: A public risk is something that could harm a lot of people, and individual potential victims have no way of stopping the harm from happening.

Examples

Nuclear technology

Environmental threats

Mass-produced consumer goods

Mechanized transportation

…Autonomous Systems?

Note: None of these really existed prior to industrialization

Industrial-Era Methods of Public Risk Management

	Formal	Informal
Preemptive	Legislation Agency rulemaking Subsidies	Industry standards
Reactive	Common law	Free market (consumer choice)

The Big Question for A.I. Risk Management

Will these industrial-era methods of risk management be capable of managing the risks associated with artificial intelligence and autonomous machines?

Shortcomings of Traditional Formal Regulation in Managing AI Risk

- **Machines are not people**
  - Legal systems operate by assigning and allocating legal rights/responsibilities to “persons” (even for corporate “persons,” it is assumed that humans make all important decisions)
  - The idea that something other than a human can make a legally significant decision is foreign to our laws
- **Foreseeability concerns**
  - Law hesitates to punish people for harm they couldn’t have foreseen
  - With machine learning, even designers may not fully understand why system does something
  - Makes it difficult to assign and allocate responsibility in a way that makes deterrence effective
- **Control concerns**
  - Autonomous systems’ priorities and incentives may not align with ours—even if we program them
- **“Wind shear”: Coping with simultaneous *atomization* and *concentration***

Atomization

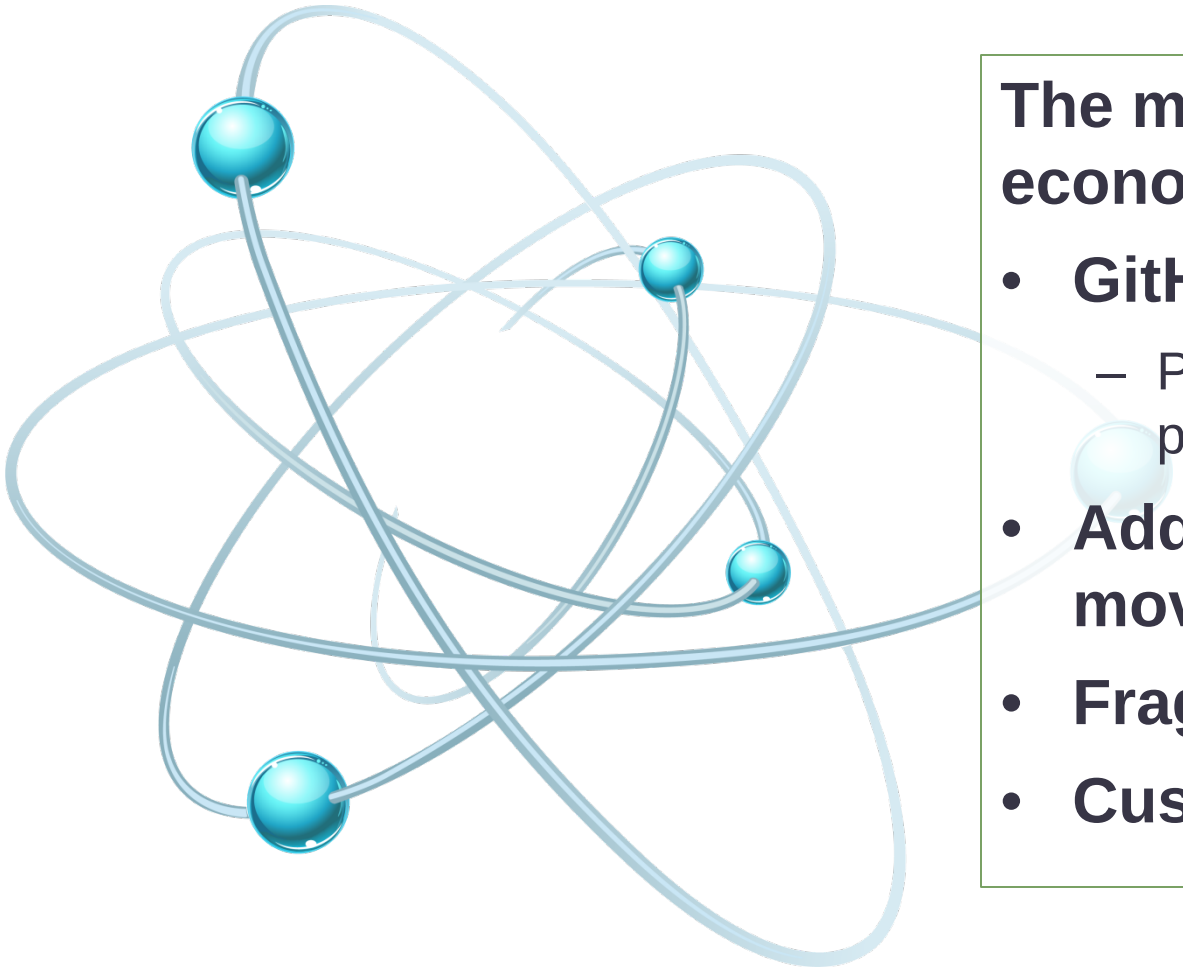

The modern world is making decentralized economic activity ever-easier

- **GitHub**
  - People all over the world can collaborate on programming projects
- **Additive manufacturing / maker movement**
- **Fragmentation of news sources**
- **Customization and personalization**

Problematic Features of Digital-Age Development

- **Discreetness**
  - Risky AI development might be done in locations and using methods that escape detection by regulators
- **Discreteness**
  - Risks might stem from the interaction of components created at different places and times, without conscious coordination
- **Diffuseness**
  - Designers and manufacturers of components may be in different jurisdictions (and operators in yet other jurisdictions)
- **Opacity**
  - Regulators may not be able to discover or understand the underlying mechanisms that create risks

Concentration

- **Revenue of Big Five tech companies in 2016: $556b**
  - (Argentina GDP: $546b)
- **They will have access to data that, in some cases, have a level of detail far exceeding that of the governments charged with regulating them**
- **Perhaps less problematic than the decentralizing forces**
  - Regulatory models were in part built to provide a counterweight to over-powerful companies
  - Could conceivably “scale up” regulatory machinery to handle new corporate superpowers

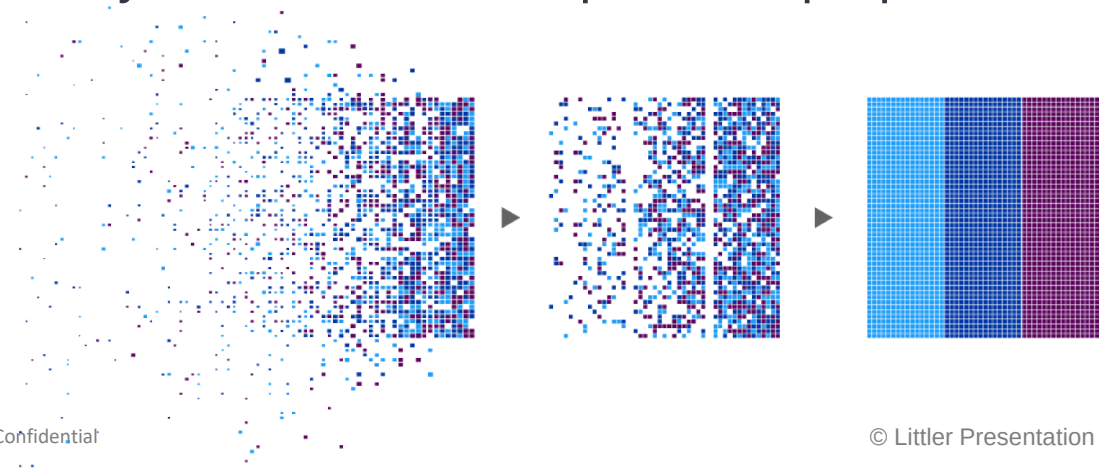

Shortcomings of the free market

- **Information asymmetry**
  - Producers have more information about risk than consumers (or regulators or competitors)
  - Particularly acute with emerging technologies
    - Even more so with Big Five
  - Failure of free market in managing industrial era risks is what led to rise of regulatory state
- **Insurance?**
  - Difficult to estimate risks with new technologies
  - Difficult to insure against large-scale public risks

Shortcomings of industry standards and self-regulation

- **Fox guarding the henhouse**
  - Industry effectively decides acceptable level of risk for public
  - Only works if industry's interests are very closely aligned with public at large.
    - Rarely the case for large companies, which traditionally generate most public risks.
- **Enforcement**
  - Market participants can avoid restrictions by simply leaving (or never joining)

Institutional Competencies: Legislatures

- **Democratic legitimacy**
  - Have best claim to be representing the interests of society at large
  - Are only institution capable of credibly establishing policy
- **Lack of expertise**
  - Inherently generalists; their ambit includes our entire economy and society
  - Typically must rely on committee hearings and contact with lobbying groups to gain access to relevant expert opinions regarding proposed legislation
  - Have committees that theoretically could allow some development of expertise, but power of committee is waning and effectiveness of committee hearings is debatable
- **Ability to delegate**
  - Accompanied by power of oversight

Institutional Competencies: Agencies

- **Specialization/Expertise**
  - Specialization: They focus all time and resources on a single industry or problem
  - Expertise: Can be staffed by technocrats and people with extensive relevant experience
  - But this edge is significantly dulled in the context of emerging technologies
- **Flexibility in structure**
  - Structure of legislatures and courts are largely static; new agencies can be designed with a structure catered to the particular industry or problem the agency is tasked with addressing
- **Independence (and Alienation)**
  - Insulated from political pressures that legislatures and even courts face
  - But can be out-of-touch or become too cozy with those who they are supposed to regulate
- **Ex ante action**
  - Legislatures rarely can react quickly enough to respond to rapidly developing crises

Institutional Competencies: Courts

- **Specialty is fact-finding and adjudication**
  - Makes courts particularly ill-suited for making findings regarding what *usually* happens in a *class* of cases, but ideally suited for finding what *actually* happened in *one specific* case.
- **Reactive (and Reactionary)**
  - Have limited power for *ex ante* action
  - Tend to treat new and unfamiliar risks far more harshly than familiar risks
- **Incrementalism**
  - Legal rules are allowed to develop slowly and organically over time; less risk of overreaction
- **Misaligned incentives**
  - Plaintiffs' lawyers choose cases based on the probability of obtaining a lucrative settlement or a favorable verdict, rather than on how best to optimize public risks
  - Lawyers focus on achieving victory in case, not providing court with info needed to make good law
  - All too easy to find "expert" witnesses who swear to something wacky

How do we regulate things we don't understand?

- **Some modern machine learning methods work in a way that makes it impossible to reverse engineer the system's "reasoning" or determine how it reached its current state**
  - Put another way, not even the people who create such systems may be able to explain its actions
- **Sounds scary but, in reality, this is not a new problem**
- **Case in point: Pharmaceuticals**
  - Throughout the history of scientific medicine, we have often recognized, tested, and adopted effective methods of treatment without understanding *why* the treatment is effective
  - Example: Smallpox vaccine was discovered at a time when we didn't know that disease was caused by germs
  - Modern Example: We still don't fully understand the mechanisms of many psychiatric drugs
- **How did we manage that risk?**
  - Heavy, regimented regulation. A product had to undergo rigorous testing and be proven (reasonably) safe before it could be marketed.

Potential New Paradigm: Crowdsourcing Regulation

- **General idea: Require transparency, then rely on stakeholders in the public at large to bring potential risks to the attention of government**
  - Inspiration: the EU's REACH regulations for the chemicals industry
- **Transparency**
  - Not in the sense of "be able to explain why a machine does what it does" (which is not always possible). Rather, "disclose enough relevant information to allow for a meaningful risk audit."
  - IP and security concerns with this approach, but that may be the cost paid to avoid public risk
- **Crowdsourced regulation**
  - Allows users, competitors, and members of general public to report potential risks
  - By making relevant details of AI systems available to everyone, chances of risk detection are maximized

Littler®

THANK YOU

Matthew U. Scherer, J.D.