Economic Implications of False Causal Attribution:

A Placebo Trilogy

Ran Spiegler (TAU & UCL)

Helsinki, December 2017

Introduction

- **Theme**: People attach value to intrinsically worthless actions because they draw false causal inferences from data.
- Today: Three explorations of this theme
  - Search for good ways to model false causal attribution
  - Different contexts: Individual behavior, markets, public policy

Introduction

- Three parables, each based on a different paper (in non-chronological order):

1. "Bayesian Networks and Boundedly Rational Expectations" (QJE 2016): **Individual behavior**
2. "The Market for Quacks" (Restud 2006): **Markets**
3. "Placebo Reforms" (AER 2013): **Public policy**

Part I:

The Dieter’s Dilemma

The Dieter's Dilemma

- An agent chooses whether to consume a dietary supplement.
- Three variables that take the values $0,1$:
  - a represents the agent's action (1 means consuming)
  - h represents state of health (1 means good health)
  - c represents blood chemical level (1 means abnormal)
- The agent's payoff is $h - ka$, where $k > 0$ is constant.

The Dieter's Dilemma

- p is an objective long-run joint distribution over a, h, c.
- $p(h = 1) = 0.5$, independently of a.
  - The objectively optimal choice is $a = 0$.
- $p(c = 1 \,| a, h) = (1 - a)(1 - h)$
  - The agent's chemical level is abnormal only if he is unhealthy **and** doesn't take the supplement.

The Dieter's Dilemma

- The agent perceives p with a subjective causal model.
- The model, denoted R, is represented by a directed acyclic graph (DAG) $a \rightarrow c \rightarrow h$.
- The agent fits his causal model to the long-run distribution:

$$p_R(a, h, c) = p(a)p(c|a)p(h|c)$$

- General belief distortion rule: Factorizing p according to a DAG

The Dieter's Dilemma

- The agent relies on p_R to evaluate the consequences of actions:

$$p_R(h|a) = \sum_c p(c|a)p(h|c)$$

- Why doesn't the agent directly estimate $p(h|a)$?
  - The model alerts the agent to some correlations, not to others
  - Data not easily available; using c as a surrogate variable.

The Dieter's Dilemma

- $a \rightarrow c \rightarrow h$ assumes $h \perp a \mid c$. This assumption is false:
  - Conditional on a normal chemical level, learning that the agent didn't take the supplement leads us to infer that he is healthy.
- $p(a, h, c) = p(a)p(h)p(c|a, h)$ is consistent with a "true" causal model $a \rightarrow c \leftarrow h$.
  - R exhibits reverse causality w.r.t the true model.

The Dieter's Dilemma

- The agent's subjective expected utility from a:

$$\sum_h p_R(h|a)(h - ka) = p_R(h = 1|a) - ka = \sum_c p(c|a)p(h = 1|c) - ka$$

- $p(h|c)$ is **not invariant** to $p(a)$.
- Choosing a to maximize this expression is well-defined only if

  we know $p(a)$.

The Dieter's Dilemma

- This observation suggests a need for an **equilibrium** notion of subjective maximization!
- "Personal equilibrium" : If $p(a) > 0$, then a maximizes subjective expected utility.
  - Need to introduce "trembles" for definition to be precise.

The Dieter's Dilemma

- Look for a personal equilibrium with full support: The agent is indifferent between the two actions.

$$p_R(h = 1|a = 1) - p_R(h = 1|a = 0) = k$$

$$\sum_c [p(c\,|a = 1) - p(c\,|a = 0)]p(h = 1\,|c) = k$$

- We only need to calculate the conditional probabilities.

The Dieter's Dilemma

- $p(c = 0|a = 1) = 1$ $\qquad$ $p(c = 0|a = 0) = 0.5$

- $p(h = 1|c = 1) = 0$ $\qquad$ $p(h = 1\ |c = 0) = \frac{0.5}{0.5+0.5p(a=1)}$

- $p(a = 1) = (1 - 2k)/2k$

- A unique interior equilibrium exists for $k \in (0.25, 0.5)$.
  - Why? Because higher long-run consumption lowers the supplement's perceived benefit.

The Dieter's Dilemma: Summary

- The agent "mistakes correlation for causation": Misreading $c - h$ correlation as a causal effect of c on h, he ends up attaching value to the worthless supplement.
- Individual maximization as an equilibrium concept
- General modeling framework: Agents factorize the objective distribution according to their DAG, and best-reply in personal equilibrium to the derived conditional probability distribution.

Part II:

A Market for False Prophets

Naïve Extrapolation from Small Samples

- In the Dieter's Dilemma, the agent "estimated" a misspecified model, using an infinitely large sample.
- Suppose the agent's DAG is $a \rightarrow h$. In reality, $h \perp a$.
- $p_R(h|a) = p(h|a)$ – consistent with rational expectations
- Repeated measurement protects the agent from this particular causal misattribution.

Naïve Extrapolation from Small Samples

- But what if the agent had a finite sample?
  - A small sample could lead to spurious $a - h$ correlation.
  - Naïve extrapolation could lead the agent (with some probability) to attach value to the worthless product.
- Echoes of the "law of small numbers" (Tversky-Kahneman 1971)
- Market implications?

A Market for False Prophets

- n profit-maximizing firms with no special forecasting skills (costlessly) guess the outcomes of horse races.
- There is a set of horses $M = \{1, \ldots, m\}$, each with an equal chance to win any given race $c \in C = [0,1]$.
- A consumer values successful prediction at 1.
- His outside option: A free, unprofessional uniform forecast

A Market for False Prophets

- The firms play a simultaneous-move game. A firm's strategy is a pair (p, f), where:
  - $p \in [0,1]$ is the price ("forecasting fee").
  - $f: C \rightarrow M$ is a forecasting strategy.
- A consumer with rational expectations would realize this is a "market for false prohpets" and opt out.

A Market for False Prophets

- Real-life markets that this parable approximates:
  - Alternative medicine
  - Charlatans in other walks of life
  - Active money management (in light of the Efficient Markets Hypothesis)

A Market for False Prophets

- The consumer evaluates forecasters according to a sampling procedure (variant on Osborne-Rubinstein 1998):
  - He recalls **one** race c, uniformly drawn from C.
  - He checks which forecasters (himself included) predicted the winner in c and chooses the cheapest among them.
  - If no one predicted the winner in c, he opts out.

A Market for False Prophets

- The consumer’s sampling procedure defines the firms’ payoff function in the game.
- Restrict attention to Nash equilibria in which firms play **deterministic** forecasting strategies.
  - Capturing a motive for systematic (yet entirely spurious) product differentiation

Equilibrium when $n \geq 2m$

- At least twice as many professional forecasters as horses
- Firms charge the competitive price $p = 0$, and each horse is predicted by at least two firms in almost every race.
- No welfare loss for consumers
- Non-competitive symmetric Nash equilibrium if we allow mixed prediction strategies.

Equilibrium when $n \leq m$

- An active market for false prophets: Firms charge the monopoly price $p = 1$ and predict distinct winners in almost every race.

- Industry profits $=$ Consumers' welfare loss $= \frac{n}{m} \cdot \frac{m-n}{m}$
  - The probability that only a professional forecaster predicted the winner in the sampled race

Equilibrium when $m < n < 2m$

- No pure-strategy equilibrium
- Restrict attention to equilibria in which firms' mixed pricing strategies are continuous.
- Lemma: In almost every race, every horse is predicted by exactly one or two firms.
  - Maximally differentiated predictions

Logic behind Product-Differentiation Lemma

- Suppose there are races in which some horse that is predicted by more than two firms.
- This class of races contains a race c with the maximal number of horses that are predicted by a single forecaster.
- Firms $1,2,3$ predict horse h in c, while firm 4 exclusively predicts h'.
- Why doesn't firm 1 switch to h' in c? It must be that firm 4 charges (probabilistically) lower prices than firm 2.

Logic behind Product-Differentiation Lemma

- Because $n < 2m$, firms have some market power: There must be a race c' in which firm 2 exclusively predicts some horse h''.
- If firm 4 does not make an exclusive prediction in c', a firm that makes the same prediction as firm 4 would like to switch to h''.
- Thus, firm 4 makes an exclusive prediction whenever firm 2 does.
- But this contradicts the definition of the race c.

Equilibrium when $m < n < 2m$

- $\mu = (2m - n)/m$ is the fraction of horses (in almost) every race) that are predicted by a single firm.
- Firms mix over prices according to the *cdf* G over $[\mu, 1]$*:*

$$G(p) = (p - \mu)/(p - p\mu)$$

- Industry profits = consumers' welfare loss $= \frac{n}{m} \cdot \frac{m-1}{m} \cdot \mu$
  - Increasing in m due to weaker competitive pressures

The Market for False Prophets: Summary

- When consumers naively extrapolate causal effects of actions from small samples, a market for worthless products can thrive.
- Lower (spurious) success rate raises equilibrium prices, product differentiation and consumer welfare loss (= industry profits).
- Implications for the money management industry?

Part III:

Placebo Reforms

Salient vs. Non-Salient Actions

- In the examples we examined so far, belief in a causal link between an action and a consequence did not discriminate between different actions.
- In many real-life situations, the tendency to attribute an outcome to an action can vary with the action's salience.

Banerjee & Duflo: “Poor Economics” (2011)

“Because most diseases that prompt visits to the doctor are self-regulating…there is a good chance that patients will feel better after a single shot of antibiotics. This naturally encourages spurious causal associations: Even if the antibiotics did nothing to cure the ailment, it is normal to attribute any improvement to them. By contrast, it is not natural to attribute causal force to inaction: If a person with the flu goes to the doctor, and the doctor does nothing, and the patient then feels better, the patient will correctly infer that it was not the doctor who was responsible for the cure.”

A Dynamic Model of Strategic Reforms

- Rational policy makers (PMs) want the public to attribute good changes to their own actions, and bad changes to their predecessors or successors.
- The public employs a naïve causal attribution rule: Changes are traced to the most recent intervention.
- Implications for strategic reform decisions (timing, risk)

A Dynamic Model of Strategic Reforms

- A sequence of PMs choose actions which may affect the continuation of a stochastic process governing x^t ($t = 0,1,\ldots$)
- The process can be interpreted as “growth”.
- In each period $t > 0$, a distinct PM observes the entire history (including x^t) and chooses an action $a^t \in A$.
- The action set A contains at least two actions, including a default action denoted 0.

Policy Makers' Payoffs

- $r(t)$ denotes the earliest $s > t$ for which $a^s \neq 0$.
- PM t's payoff is 0 if $a^t = 0$, and

$$(1-\delta)\sum_{s=t+1}^{\infty} \delta^{s-t-1} x^{\min\{s,r(t)\}} - x^t$$

when $a^t \neq 0$, where $\delta \in (0,1)$ is a discount factor.

  - When δ is high, the payoff is approximately $x^{r(t)} - x^t$.

A Dynamic Model of Strategic Reforms

- Interpretation: People perceive salient action as causes.
  - A recent, non-default action is salient.
- PMs get no credit if they choose the default.
- If a PM intervenes ($a \neq 0$), all changes until the next intervention are attributed to his own intervention.
- A high δ means that PMs care about "posterity".

The Stochastic Process

- Every $a \neq 0$ is associated with a symmetric, everywhere-positive density f_a (F_a denotes the induced *cdf*).
- $x^t - x^{t-1} = \mu + \varepsilon^t - \varepsilon^{t-1}$:
  - $\mu > 0$ is a constant, action-independent trend.
  - ε^t is an independent draw from $f_{p(t)}$, where $p(t)$ denotes the latest $s < t$ for which $a^s \neq 0$.
- All interventions are "placebo reforms".

Subgame Perfect Equilibrium

Proposition: In subgame perfect equilibrium, each PM t chooses:

1. $a = 0$ if $\varepsilon^t > \varepsilon^*$.
2. $a^* \neq 0$ that minimizes $\int_{z \leq \varepsilon^*} F_a(z)$ if $\varepsilon^t < \varepsilon^*$.

where ε^* is uniquely defined by the equation

$$\mu - \delta \int_{z \leq \varepsilon^*} F_{a^*}(z) = (1 - \delta)\, \varepsilon^*$$

Equilibrium Properties

- Reform timing follows a stationary cutoff rule.
  - PMs tend to intervene after bad shocks and get undeserved credit for subsequent changes.
- Risk aversion in choice of interventions
  - $\int_{z \leq \varepsilon^*} F_a(z)$ measures the riskiness of a.
  - ε^* decreases with riskiness of a^*.

A Strategic Multiplier

- A non-strategic variant: no one moves after PM t.
  - The PM intervenes only if $\varepsilon^t < \varepsilon^{**}$, where $\mu = (1-\delta)\varepsilon^{**}$.
- $\varepsilon^* < \varepsilon^{**}$: Anticipation of future PMs' strategic reforms exacerbates the tendency to intervene during crises.
- Risk aversion disappears in the non-strategic case.

Placebo Reforms: Summary

- When the public attributes outcomes to recent interventions, PMs tend to act during crises.
- Anticipation of successors' strategic reform decisions exacerbates the PM's tendency to act during crises, and adds a risk-aversion motive.
- Bad effects on growth when actions aren't pure placebos

Conclusion

- Why am I so obsessed with the "placebo" theme?
- Amateur psychoanalysis: Unease with some people/activities that our society tends to reward (star CEOs, hedge fund managers, education gurus, economics professors…)?
- A methodological reason: An extreme, and therefore illuminating example of the broader theme of how false causal attribution affects subjective valuation.

Conclusion

- This broader theme has far-ranging economic implications:
  - Markets for credence goods (pricing, product complexity/differentiation, disclosure, comparative statics)
  - Money management
  - Speculative trade in financial markets
  - Education, health and lifestyle decisions by individuals
  - Macroeconomic policy with a boundedly rational private sector
- The DAG formalism as a basis for a unifying framework