“Synthetic Control As Online Linear Regression” by Jiafeng Chen

Jann Spiess, Stanford

Program evaluation meets online learning

Panel econometrics | Online convex optimization

learn on past → Take accurate decision now

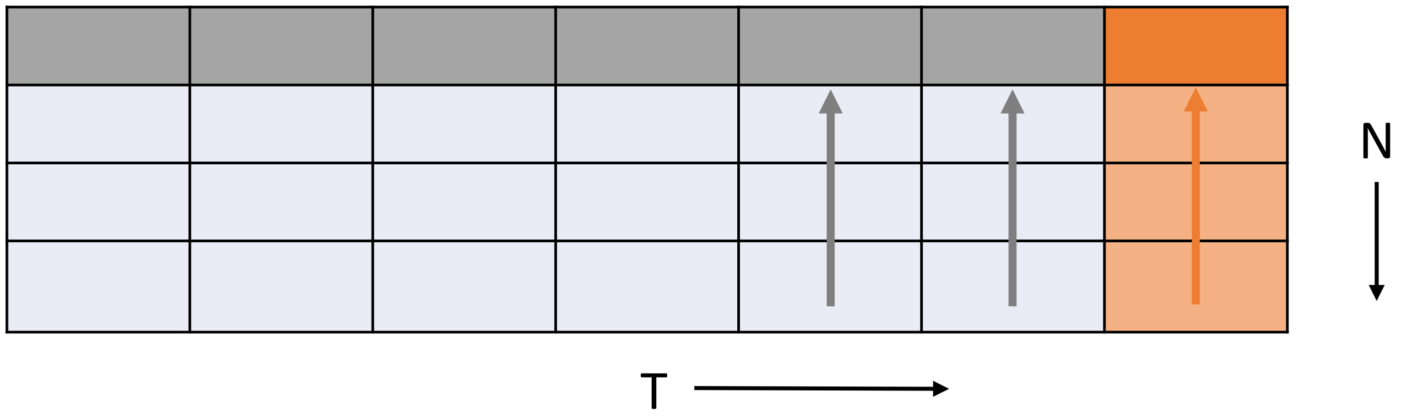

Synthetic Control

Follow-the-L

- Similar structure and goal
- Turns out that solution also takes similar form
- **So what is different? What does it add? Why should we care?**

Different goals, different results

	Applied Econometrics	This approach
DGP	Specific model (e.g. factor model)	Worst case chosen by adversary
Estimator	Ad-hoc/derived from model (e.g. maximum likelihood)	Class of estimators (e.g. convex combinations)
Criterion	Risk/MSE, variance subject to unbiasedness	Regret relative to oracle (e.g. $E[\left(y_{0\tau} - \hat{y}_{0\tau}(\hat{\theta}_\tau)\right)^2] - \min_{\theta \in \Theta} E[\left(y_{0\tau} - \hat{y}_{0\tau}(\theta)\right)^2]$)
Result	Consistency, (asymptotic) distribution	Regret rates (e.g. $O(\log(T)/T)$)

Specific assumptions and criteria for synth

- <u>DGP</u>: Outcomes can be anything, only assume bounded
- <u>Accuracy</u>: Expected/average squared error $E[\left(y_{0\tau} - \hat{y}_{0\tau}(\hat{\theta}_\tau)\right)^2]$
- <u>Estimator class</u>: Only consider convex averages
  $\hat{y}_{0t}(\theta) = \theta' \boldsymbol{y_t} = \sum_i \theta_i y_{it}, \theta \in \Theta = \{\theta; \theta_i \geq 0, \sum_i \theta_i = 1\}$
- <u>Reference estimator</u>: Best (oracle) estimator given actual outcomes y_{it} that chooses same parameter θ every period, $\underset{\theta \in \Theta}{\arg\min}\ E\left[\left(y_{0\tau} - \hat{y}_{0\tau}(\theta)\right)^2\right]$
  → important assumption: stability over time
- <u>Main criterion</u>: *feasible* estimator close to reference as $T \to \infty$ in worst case
- That means that the nature of the result is different from what we are used to:
  (It makes some specific asymptotics/invariances)
  1. It does not guarantee that synthetic control does well in absolute terms
  2. But guarantees that it is best among class under minimal assumptions

Econometrics vs online learning?

	Applied Econometrics	This approach
DGP	Factor model	Agnostic worst-case (implicit time invariance)
Estimator	Synthetic control	Any convex combination
Criterion	Risk/MSE, variance subject to unbiasedness	Regret relative to persistent oracle
Result	Consistency	Regret rate $\mathrm{O}(\log(T)\ /T)$

Econometrics and online learning!

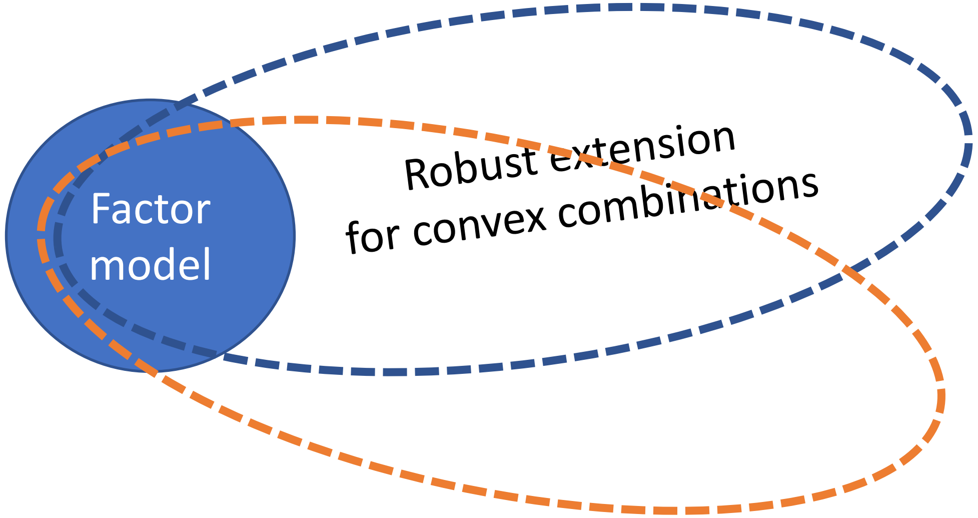

- Provides **complementary** properties of synth that draws strength from combination of econometric context with the structure of the optimization

- Really hard to derive properties of synth using standard methods!

- Suggests extensions, allows transfer from well-established framework

Extensions I am curious about

$$O(N\ \log(T)\ /T)$$

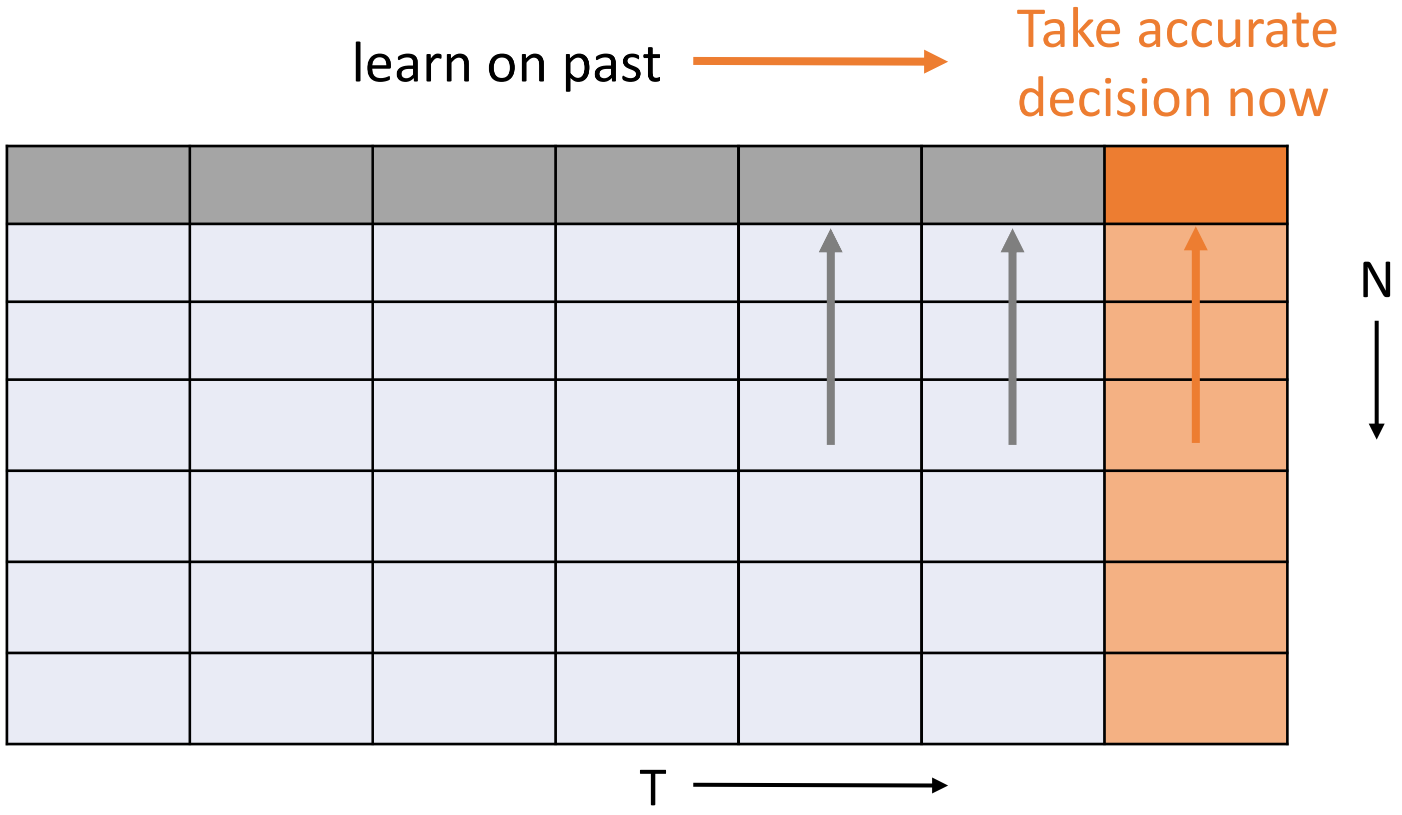

Extensions I am curious about

- What happens if we change the class of estimators?

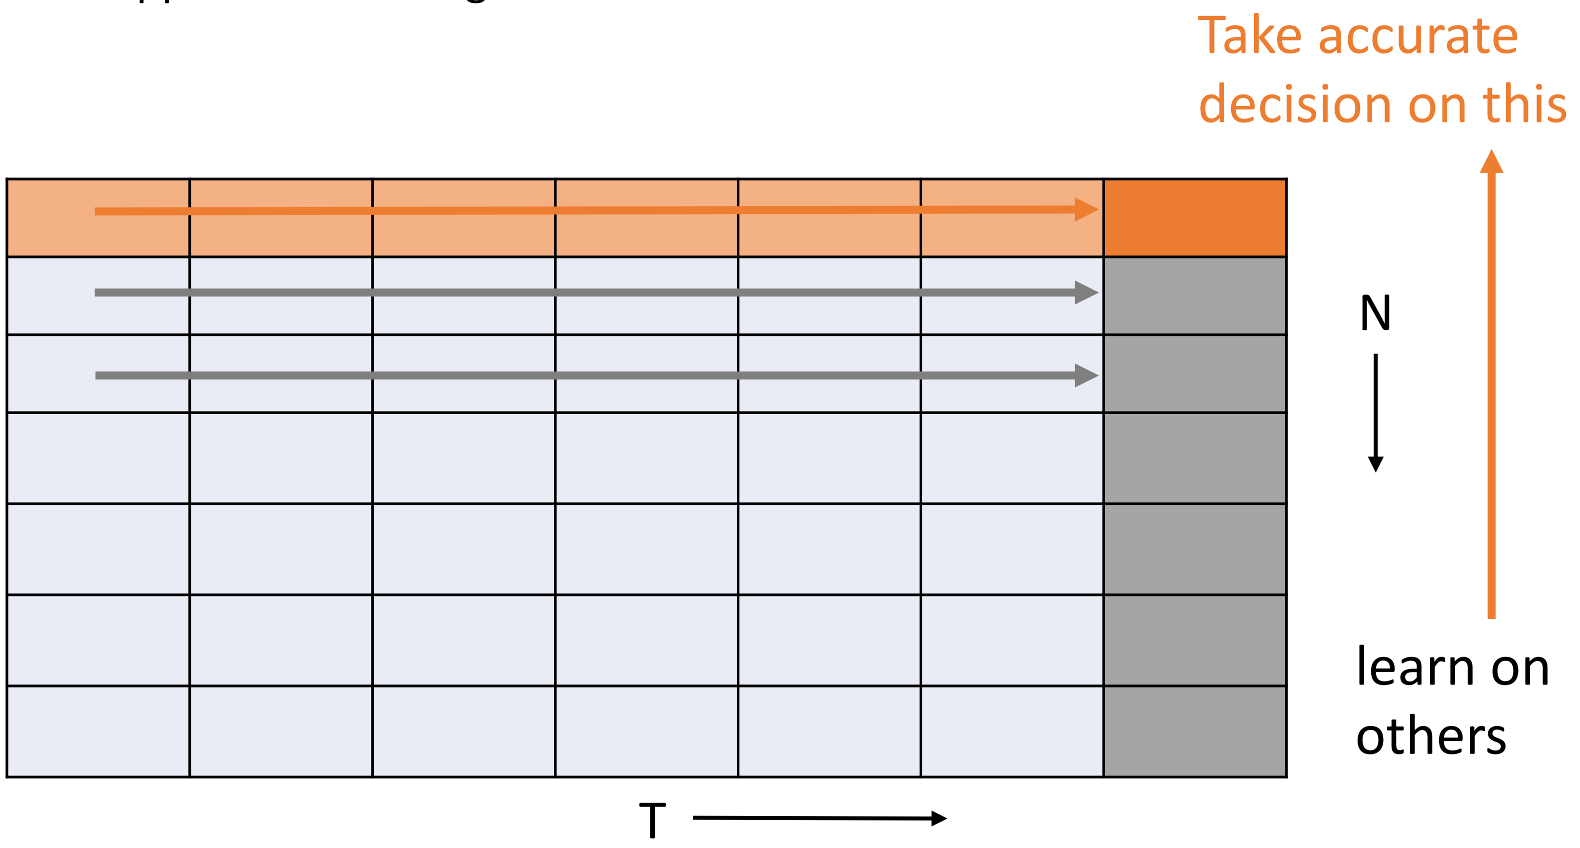

Extensions I am curious about

- What happens if we make the class of estimators larger? (matrix completion view)

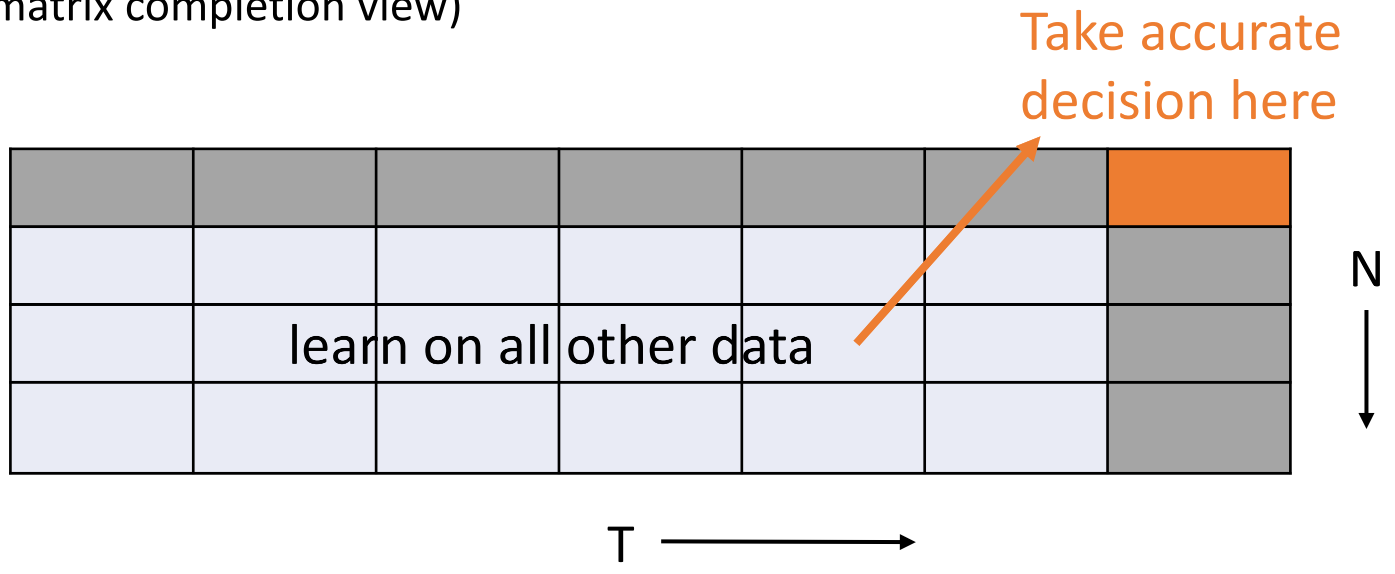

Extensions I am curious about

- What happens if we make the class of estimators larger?

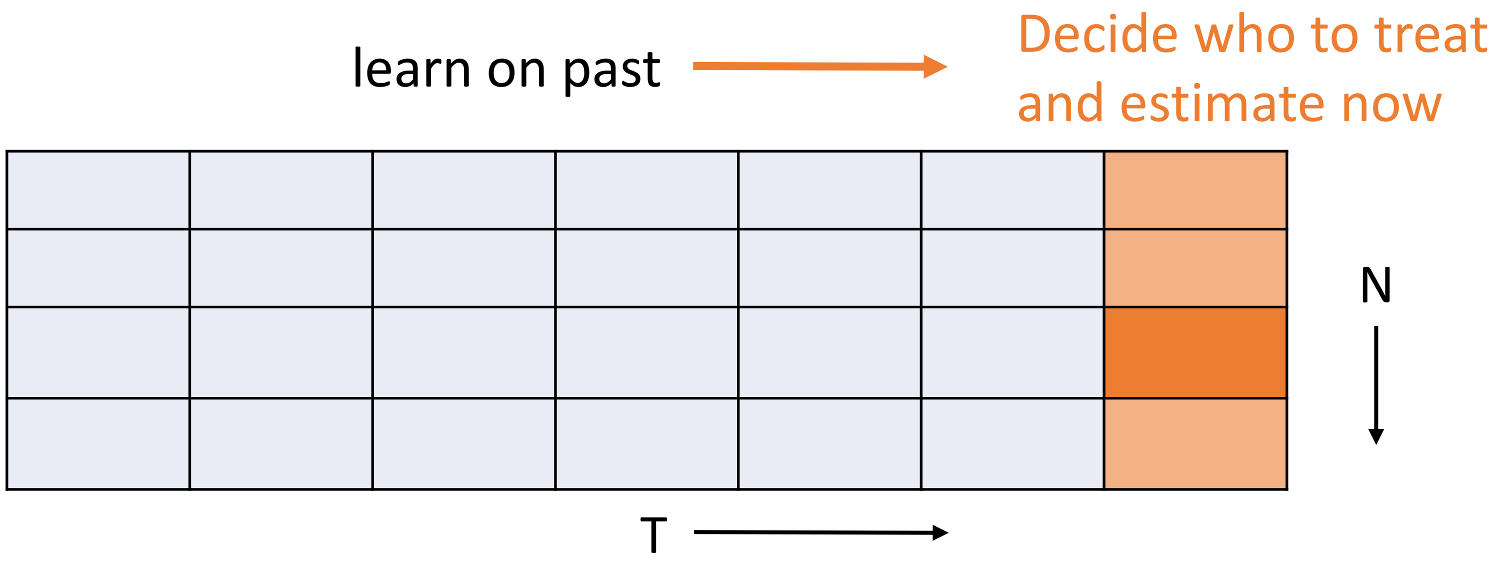

- What happens if we tweak the question?
  - Modify loss → regularized synthetic control
  - Consider dynamic experimentation (papers today; Abadie & Zhao)

Broader context and implications

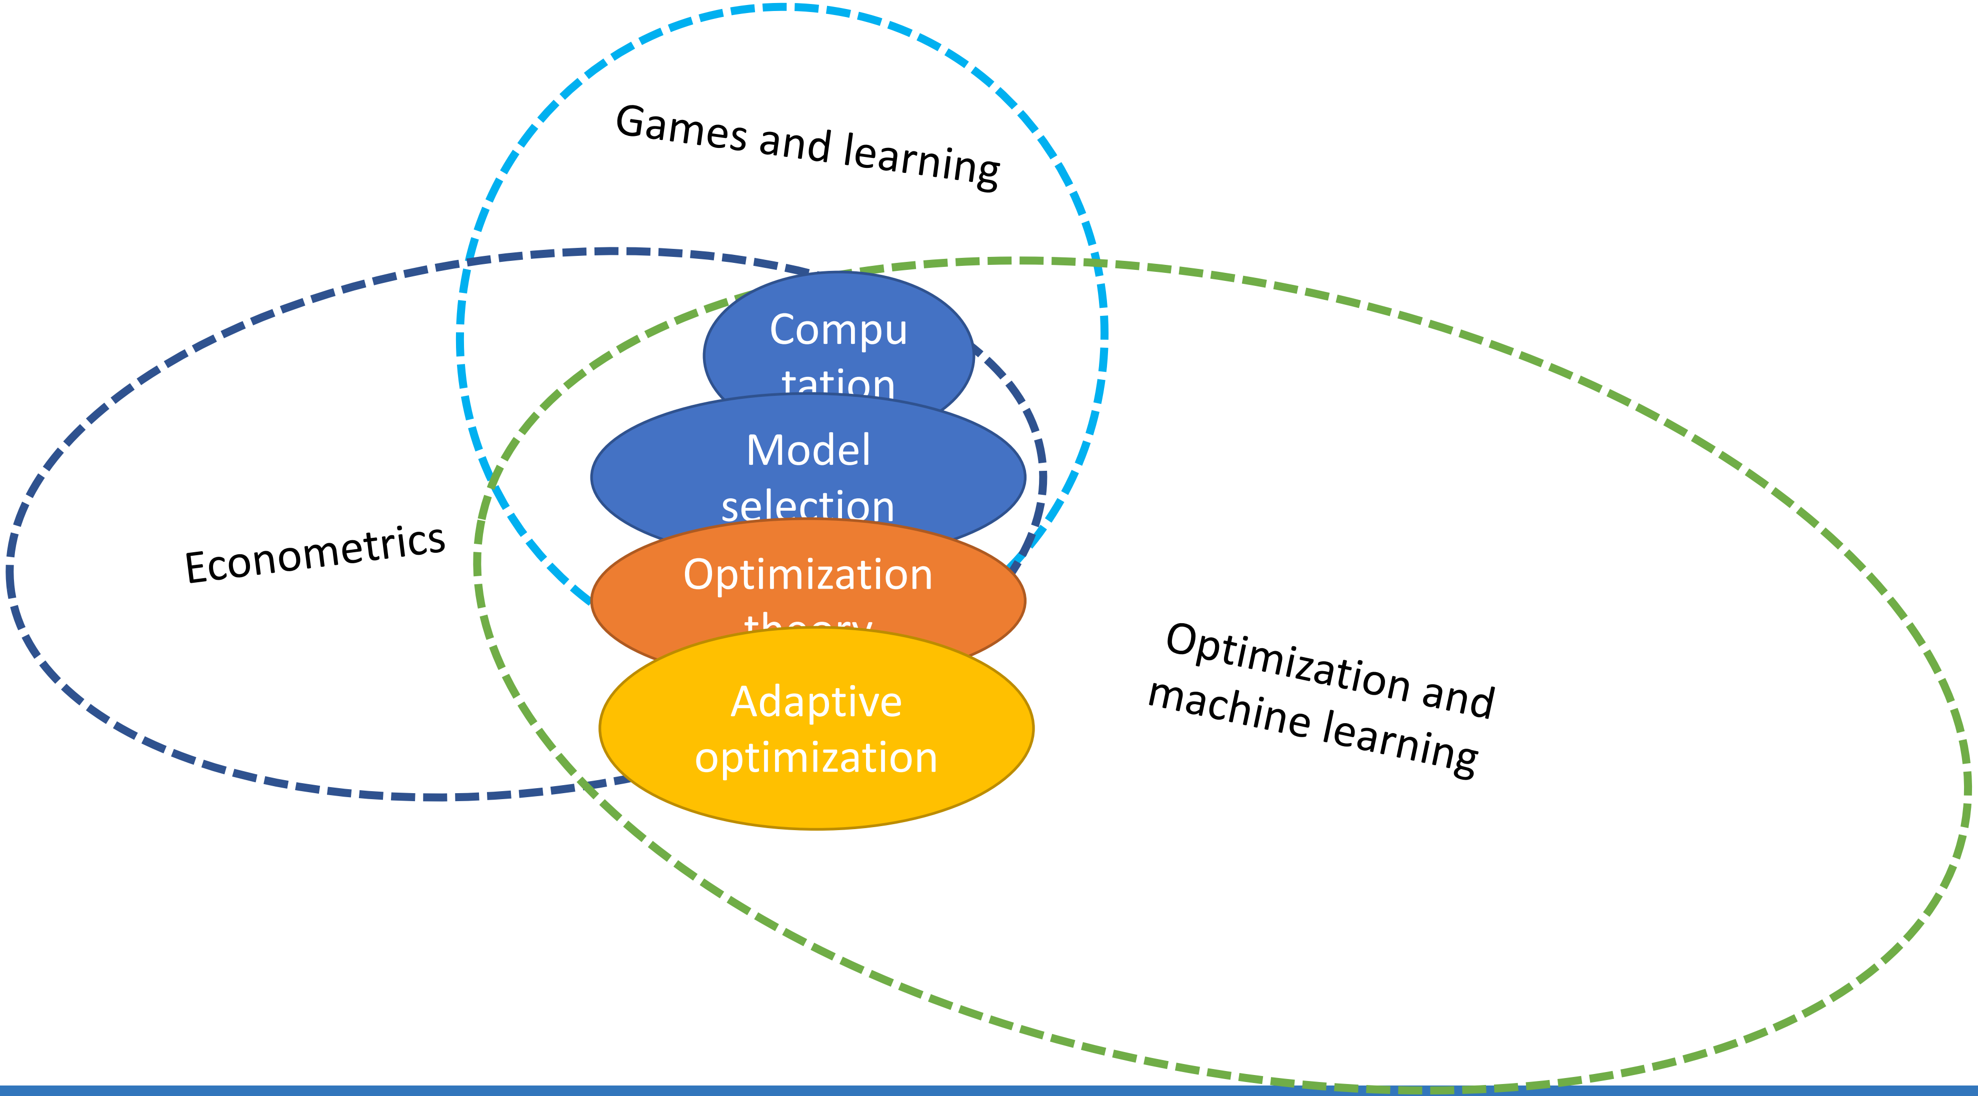

"Synthetic Control As Online Linear Regression" by Jiafeng Chen

Not just a technical connection to derive otherwise hard-to-prove properties of one (very) important estimator, also a model for optimization-based approaches for developing, choosing, and understanding the applicability of econometric methods.