Informational Substitutes and Complements

Bo Waggoner
University of Pennsylvania

June 18, 2018

Background

- **Substitutes**: the total is less valuable than the sum of individual values
  each's value decreases given some of the others
- **Complements**: the total is more valuable than the sum of values
  each's value increases given some of the others

e.g. bread and pasta; two weather channels
left shoe/right shoe; thermometer and humidity reading

- S&C for items are useful concepts in economics and algorithms;
  how do they apply to **signals**?

What we will cover

- Definitions of **informational substitutes and complements**
  [Waggoner, Chen 2016; Börgers, Hernando-Veciana, Krähmer 2013]

- Known applications and results:
  - equilibria of prediction markets
  - algorithms for S&C [Kong, Schoenebeck 2018]
  - information acquisition
    connections to experimental design, statistics, econ

- Known examples / classes of S&C

- Open problems and directions

Outline

1. Marginal value; definitions of S&C
    - relationship to information theory
    - moderate and strong definitions
2. Known applications
    - prediction markets
    - algorithmic problems
3. Known classes of examples
4. Open problems; directions

Marginal value; definitions

Reminder: setting and value function

- Prior p on states Θ
- Set of signals: $\Sigma_1, \ldots, \Sigma_n$. — σ_i is a realization of Σ_i
- Distribution $\varphi(\sigma_1 \ldots \sigma_n, \theta)$ — $\Pr[\ \ldots \Sigma_i = \sigma_i \ldots \mid \theta\ ]$
- Decision problem $u(a, \theta)$

Value function:

$V^{u,\varphi}(\Sigma)$ = "expected utility given Σ" — assuming optimal action

$= \mathbf{E}_{\sigma\sim\varphi} \max_a \mathbf{E}_{\theta|\sigma} u(a, \theta)$ — p_σ = posterior on θ given σ

$= \mathbf{E}_{\sigma\sim\varphi} G(p_\sigma)$. — G = convex expected utility function

$V^{u,\varphi}(\perp) = G(p)$ — expected utility with no information

- Extends to any subset or garbling of the "base" signals $\Sigma_1, \ldots, \Sigma_n$.

Key example: log scoring rule

Example: $u(q, \theta) = \log q(\theta)$.

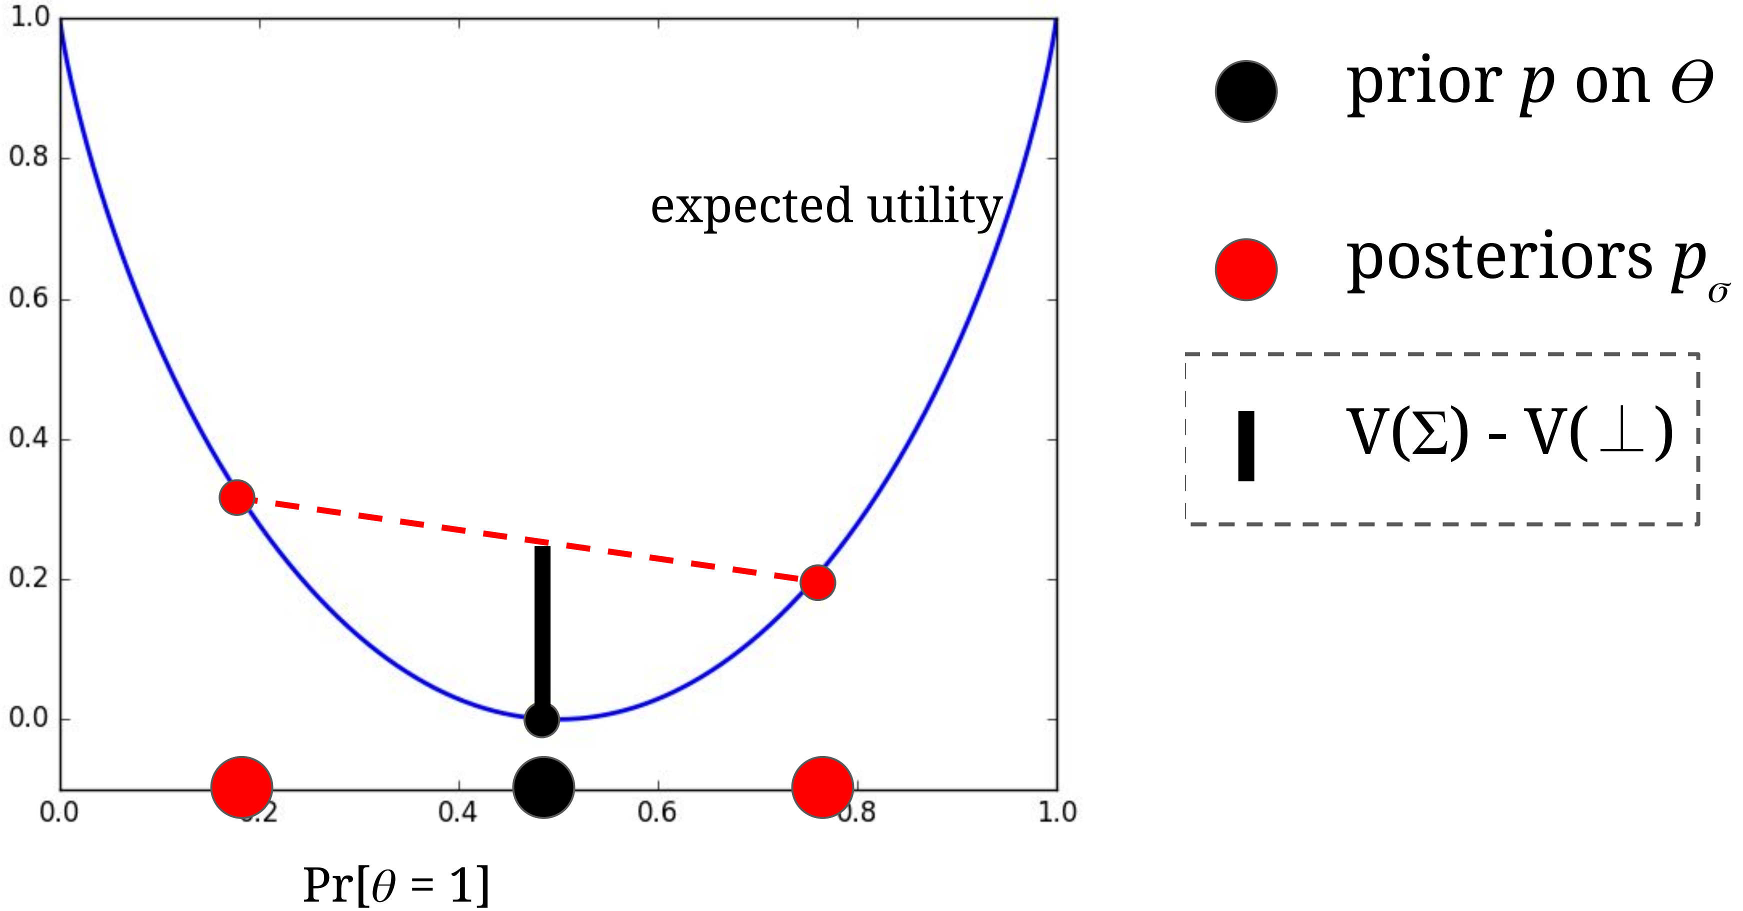

Definition: (weak) substitutes

Def. $\Sigma_1, \ldots, \Sigma_n$ are **(weak) substitutes** for u if $V^{u,\varphi}$ is *submodular*, i.e. for all Σ_i and all $S \subseteq T \subseteq \{\Sigma_1, \ldots, \Sigma_n\}$,

$$V^{u,\varphi}(S \cup \Sigma_i) - V^{u,\varphi}(S) \geq V^{u,\varphi}(T \cup \Sigma_i) - V^{u,\varphi}(T).$$

"The marginal value of Σ_i is decreasing in knowledge of the other signals."

They are **(weak) complements** for u if $V^{u,\varphi}$ is *supermodular* (ineq reversed).

Depends on both the information structure and the decision problem!

Initial examples

Canonical substitutes: for any decision problem

With probability 1, $\Sigma_1 = \ldots = \Sigma_n$.

Canonical complements: also for any decision problem

Each Σ_i i.i.d. uniform {0,1}

θ = XOR of all the bits.

Intuitively substitutes: causation: $\theta \rightarrow \Sigma_1 \ldots \Sigma_n$

relatively low sensitivity

Conditionally independent noisy observations of θ.

Intuitively complements: causation: $\Sigma_1 \ldots \Sigma_n \rightarrow \theta$

relatively high sensitivity

Independent components of a system or function θ.

Interpretations of the marginal value function

1. Value.

Given a utility function u, consider — e.g. log scoring rule

$V^{u,\varphi}(\Sigma) - V^{u,\varphi}(\perp)$ = *marginal value* of Σ.

2. Distance.

Given a Bregman divergence d, consider — e.g. KL-divergence

$\mathbf{E}_\sigma\, d(p_\sigma, p)$ = *average change in belief* due to Σ.

3. Uncertainty.

Given a concave "entropy" H, consider — e.g. Shannon entropy

$H(\Theta|\Sigma) - H(\Theta)$ = *information conveyed* by Σ. — where $H(\Theta|\Sigma) := \mathbf{E}_\sigma H(p_\sigma)$

Fact: There is a 1-1-1 correspondence between u, d, H such that

$V^{u,\varphi}(\Sigma) - V^{u,\varphi}(\perp) = \mathbf{E}_\sigma\, d(p_\sigma, p) = H(\Theta|\Sigma) - H(\Theta)$. — e.g. above examples

Marginal value = expected KL-divergence

Example: $u(q, \theta) = \log q(\theta)$.

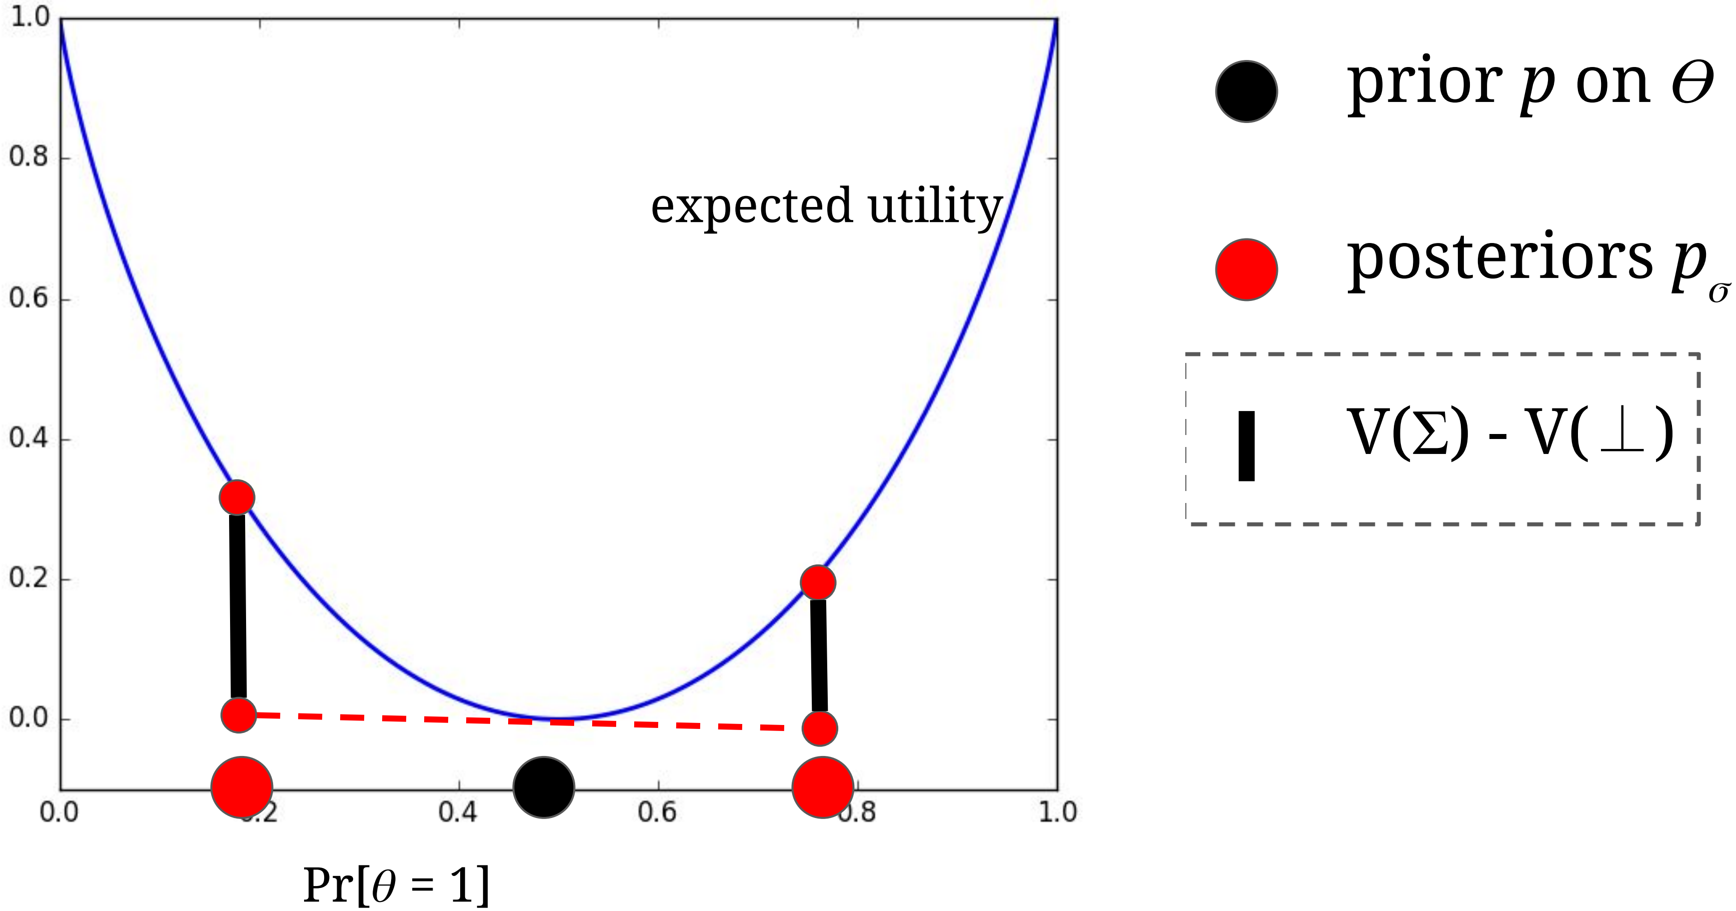

Marginal value = reduction in entropy of Θ

Example: $u(q, \theta) = \log q(\theta)$.

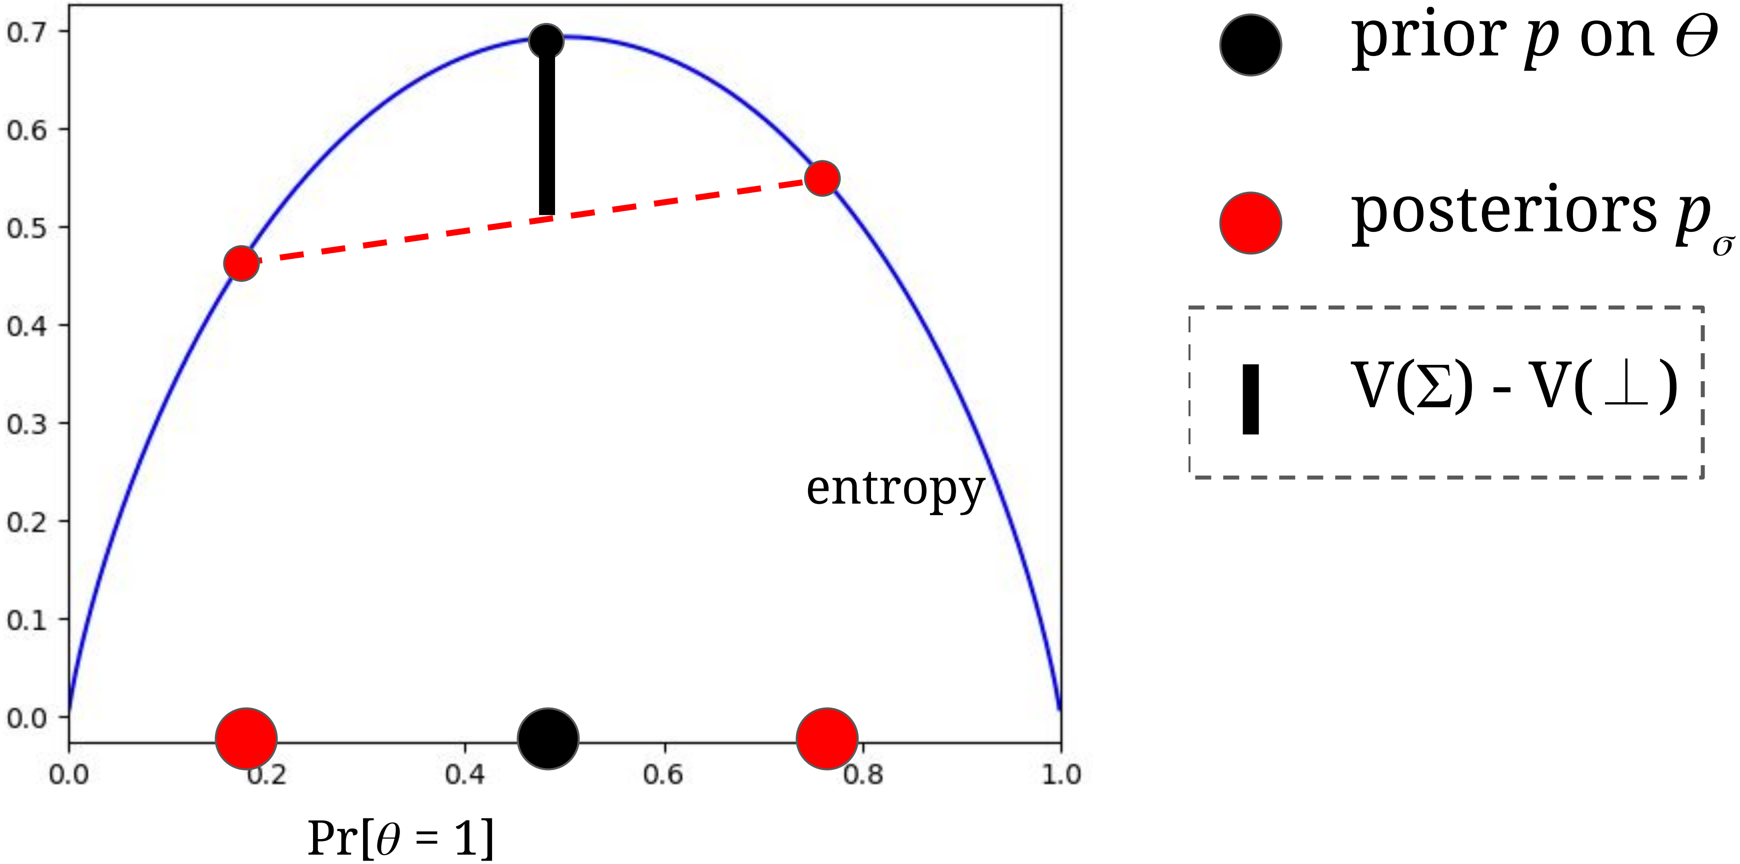

Definition: (weak) substitutes

Def. $\Sigma_1, \ldots, \Sigma_n$ are **(weak) substitutes** for *u* if $V^{u,\varphi}$ is *submodular*, i.e. for all Σ_i and all $S \subseteq T \subseteq \{\Sigma_1, \ldots, \Sigma_n\}$,

$$V^{u,\varphi}(S \cup \Sigma_i) - V^{u,\varphi}(S) \geq V^{u,\varphi}(T \cup \Sigma_i) - V^{u,\varphi}(T).$$

"The marginal value of Σ_i is decreasing in knowledge of the other signals."
"The change in belief (distance) due to Σ_i is decreasing in ...."
"The marginal amount of information contained in Σ_i is decreasing in ...."

They are **(weak) complements** for *u* if $V^{u,\varphi}$ is *supermodular* (ineq reversed).

Depends on both the information structure and the decision problem!

Drawback of “weak” defs: signals are divisibile

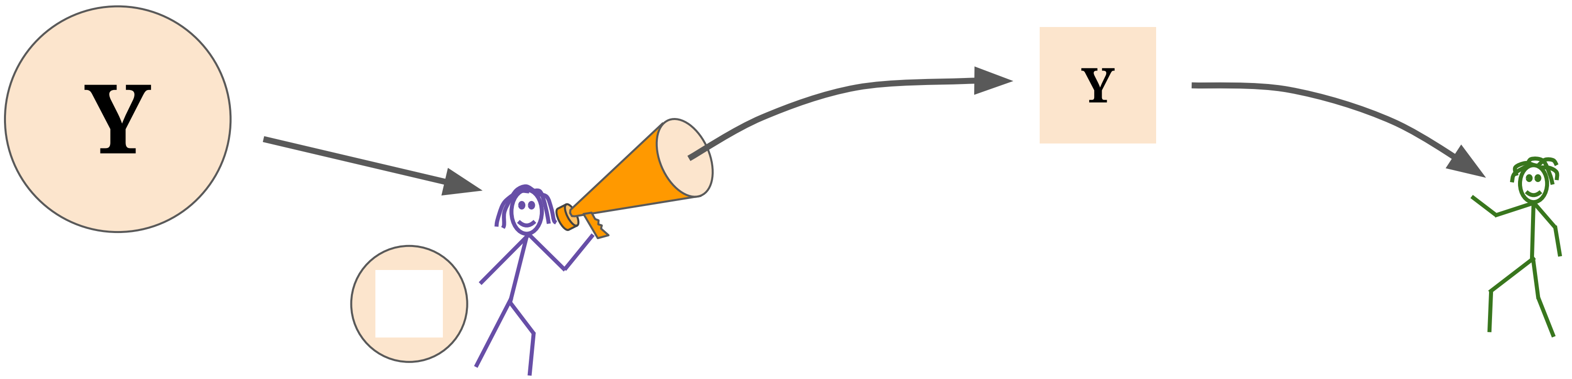

“Half the truth is often a great lie.”
- Benjamin Franklin

Example: Alice observes entire stock market,
but strategically reports one stock’s performance.

Partial revelation and stronger definitions

Given a signal Σ_i and a function f, let $f(\Sigma_i)$ be the signal consisting of $f(\sigma_i)$.

Def. $\Sigma_1, \ldots, \Sigma_n$ are **(moderate) substitutes** for u if for all **deterministic f**, for all Σ_i and all $S \subseteq T \subseteq \{\Sigma_1, \ldots, \Sigma_n\}$,

$$V^{u,\varphi}(S \cup f(\Sigma_i)) - V^{u,\varphi}(S) \geq V^{u,\varphi}(T \cup f(\Sigma_i)) - V^{u,\varphi}(T).$$

They are **(moderate) complements** for u if the inequality always reverses.

Def. $\Sigma_1, \ldots, \Sigma_n$ are **(strong) substitutes** for u if for all **randomized f**, for all Σ_i and all $S \subseteq T \subseteq \{\Sigma_1, \ldots, \Sigma_n\}$,

$$V^{u,\varphi}(S \cup f(\Sigma_i)) - V^{u,\varphi}(S) \geq V^{u,\varphi}(T \cup f(\Sigma_i)) - V^{u,\varphi}(T).$$

They are **(strong) complements** for u if the inequality always reverses.

Comment on definitions

- *Weak* definition seems uncontroversial and general.
- *Moderate* and *strong* definitions are somewhat tailored to the prediction market application.
- Future applications may need to tweak details:

  "marginal value of _____ when added to ______ must be diminishing"

Moderate:	deterministic garbling	any subset of the signals
Strong:	randomized garbling	any subset of the signals

Known applications

“Marginal-score games”

Decision problem $u(a, \theta)$. Players 1…n have private signals $\Sigma_1, \ldots, \Sigma_n$.

1. Players take turns proposing actions $a^1 \ldots a^T$. multiple plays allowed
2. θ is revealed.
3. Reward for update $a^{t-1} \rightarrow a^t$ is $u(a^t, \theta) - u(a^{t-1}, \theta)$.
4. $\Rightarrow$ player’s reward is sum of marginal improvements.

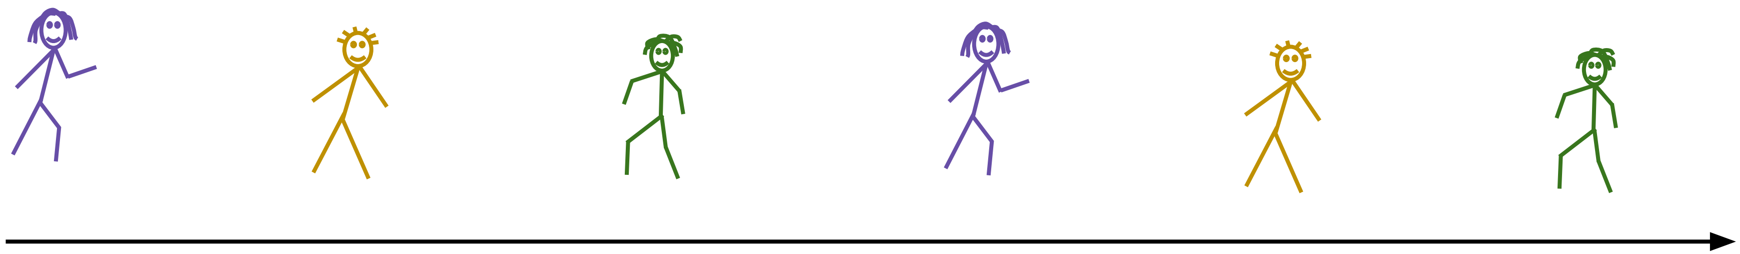

“Marginal-score games”

Decision problem $u(a, \theta)$. Players 1…n have private signals $\Sigma_1, \ldots, \Sigma_n$.

1. Players take turns proposing actions $a^1 \ldots a^T$. multiple plays allowed
2. θ is revealed.
3. Reward for update $a^{t-1} \rightarrow a^t$ is $u(a^t, \theta) - u(a^{t-1}, \theta)$.
4. $\Rightarrow$ player’s reward is sum of marginal improvements.

Theorem (Chen, Waggoner 2016):

1. The only Bayes-Nash equilibria are to play **always myopically optimally** if and only $\Sigma_1, \ldots, \Sigma_n$ are strong **substitutes** for u.

2. The only perfect Bayesian equilibria are **copy the previous strategy until your final participation, then play optimally** if and only if $\Sigma_1, \ldots, \Sigma_n$ are strong **complements** for u.

Details: assume signals are inferrable from actions; any order of participation is allowed.

Proof idea

Total utility available is $V^{u,\varphi}(\Sigma_1, \dots, \Sigma_n) - V^{u,\varphi}(\perp)$.

assume a^0 is optimal for the prior

In equilibrium, let S^t be all information revealed up to time t. Player *i* can obtain $V^{u,\varphi}(S^t \cup f(\Sigma_i)) - V^{u,\varphi}(S^t)$ for any randomized strategy f.

Substitutes $\Longleftrightarrow$ always optimal to reveal **earlier**.

Complements $\Longleftrightarrow$ always optimal to **delay**.

Therefore: every* equilibrium has the stated form.

*proof is nonconstructive for beliefs/actions off the equilibrium path! Strongly uses that "nobody is deceived in equilibrium". See also [Gao, Zhang, Chen EC'13] for a constructive example.

Examples of marginal-score games

Prediction markets.

u is a proper scoring rule and actions are predictions of θ.

[Chen, Reeves, Pennock, Hanson 2007], [Dimitrov, Sami EC'08], [Gao, Zhang, Chen EC'13], [Kong, Schoenebeck ITCS'18]

Market-based machine-learning contests.

u is a loss function, and actions are hypotheses, θ is a test data set.

[Abernethy, Frongillo NIPS'11], [Waggoner, Frongillo, Abernethy NIPS'15], [Frongillo, Waggoner ITCS'18]

Crowdsourced Q&A forums.

Decisionmaker solicits information, rewards answers proportional to value.

[Jain, Chen, Parkes EC'09]

Algorithmic problems (1)

Information acquisition problem:

Input: $\Sigma_1, \ldots, \Sigma_n$, u, and φ; prices $\pi_1 \ldots \pi_n$ for the signals; budget constraint B.

Output: subset S of signals to acquire.

Fact:

If signals are substitutes, there is a polynomial-time 1-1/e approximation.

In general or for complements, a nonzero approximation is computationally hard.

Proof: reduction to and from submodular set function maximization.

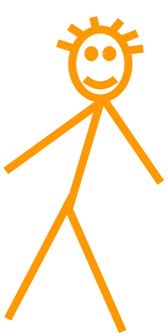

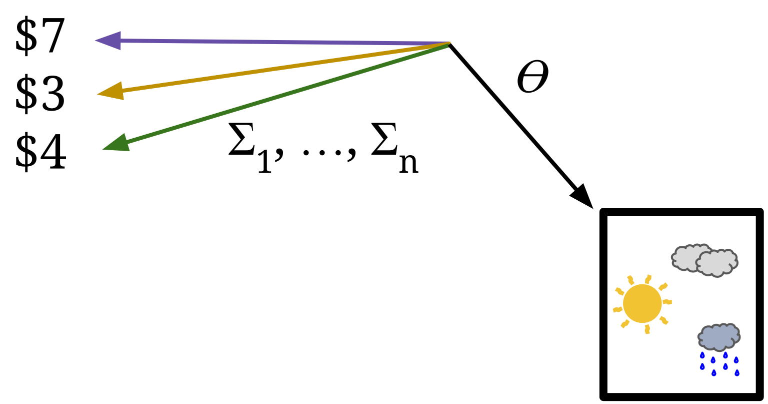

Open: more algorithmic connections

- Extends to *dynamic* information acquisition; but more to investigate.

- Point of possible connection between **games and algorithms**
    - information acquisition literature (stats, econ, CS)
    - see *Optimal and Myopic Information Acquisition* - Liang, Mu, Syrgkanis tomorrow morning -- does not use these definitions but very related!

- Econ literature models: often capture substitutes with *positive correlation*
    - connection between these?

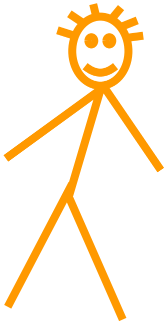

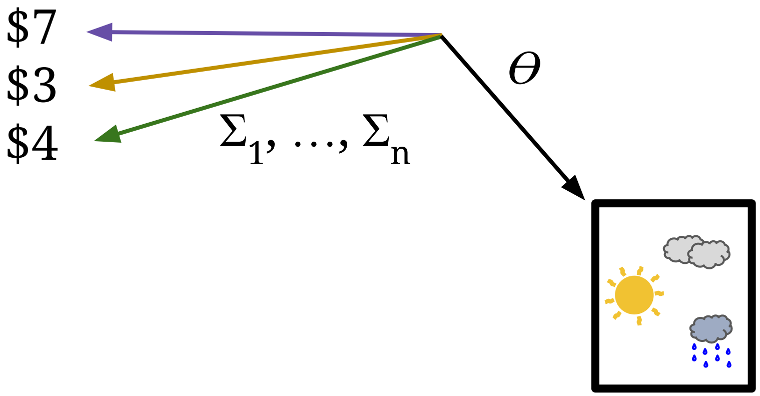

Algorithmic problems (2)

S&C identification problem:

Given $\Sigma_1, \ldots, \Sigma_n$, u, and φ:

Are they **substitutes**, **complements**, or **neither**?

exponential # of subsets; not obvious with n=2

Marginal value optimization problem:

Given Σ_1, Σ_2, compute a garbling of Σ_1 to minimize marginal value of Σ_2:

$$\text{argmin}_f \; V^{u,\varphi}(\Sigma_2 \cup f(\Sigma_1)) - V^{u,\varphi}(f(\Sigma_1)).$$

S or **C** ⇒ trivial answer. Restrict f for *weak*, *moderate* versions. Gives best-responses in prediction markets!

Theorem (Kong, Schoenebeck 2018):

There is an FPTAS for the marginal value optimization problem, treating the number of outcomes $|\Sigma_1|$ as fixed.

⇒ efficient test for identifying approximate S&C for small number of signals, and identifying all-rush or all-delay equilibria in prediction markets.

What do we know about S&C?

Knowledge about classes of S&C

For the **log scoring rule** decision problem:

- Signals *conditionally independent* on θ are strong **substitutes**
- Have separations between *weak*, *moderate*, and *strong* substitutes

  [Kong+Schoenebeck ITCS'18]

For **every decision problem** where G has a jointly convex Bregman div.:

- Signals *unconditionally independent* are strong **complements**

When are signals $\Sigma_1, \ldots, \Sigma_n \sim \varphi$ substitutes for every u?

[Börgers, Hernando-Veciana, Krähmer JET 2013]: define two signals to be substitutes if they are **weak substitutes** for **every decision problem**.

→ *universal weak substitutes* in my terminology.

→ results on structure (universality is very restrictive)

→ universal moderate/strong S&C are trivial [Anunrojwong, Chen, Waggoner]

Open problems and directions

Problems and directions

Applications:

- Financial markets; efficient market hypothesis.
  [Ostrovsky 2013]
- Common-value auctions.
  [Milgrom, Weber 1982. *The Value of Information in a Sealed-Bid Auction*]
- More general mechanism design?
  usefulness of S&C is open…

Structure and algorithms:

- Identify natural *classes* of S&C
  e.g. sets {decision problems, signals} where all combos are substitutes
- Closer connections to information acquisition
- Improve on KS18 or show hardness
- Identify conditions for efficient algorithms

Thanks!

Questions?