Signaling With Commitment

Ralph Boleslavsky Mehdi Shadmehr

NYU
April 2024

Introduction

Classic signaling game, (Fudenberg and Tirole pp. 324-5)

- Nature selects a state $\omega \in \Omega$ according to prior distribution μ_0.
- Sender observes ω, selects action $s \in S$.
- Receiver observes s but not ω, chooses $a \in A$.
- Payoffs: sender $v(\omega, s, a)$, receiver $u(\omega, s, a)$.
- Simplification: $|\Omega| < \infty$ and $|S| < \infty$.
- Applied to every topic under the sun.

Introduction

Our version endows sender with *commitment power*.

- Sender commits to his strategy, $\pi(s|\omega)$ for all $s \in S$ and $\omega \in \Omega$.
- Nature selects a state $\omega \in \Omega$ according to prior distribution μ_0.
- Sender's action is realized from his strategy.
- Receiver observes s and $\pi(\cdot|\cdot)$, but not ω. Chooses $a \in A$.
- Payoffs: sender $v(\omega, s, a)$, receiver $u(\omega, s, a)$.
- Commitment power can come from design of institutions, formal contracts, reputation incentives, algorithms/AI.
- Alternate Interpretation: sender commits to a statistical experiment on state, payoffs depend directly on *realization* of experiment

Contribution

Theory

- Study problem using the "belief-based" approach
- Geometric characterization of sender's attainable payoffs
- Characterize "extended commitment": communication protocol and action

Applications

- Rating Design: "exploiting credulity" vs. "costly lies"
- Platform Design: "steering" vs. "information provision"

Example

Adjudication.

- Grievances arise in organization, valid $\omega = v$, or invalid $\omega = f$
- Prior belief: $\mu_0 = \Pr(\omega = f)$, generally μ is prob. of f
- Org. (sender) observes type of grievance, $\omega \in \{v, f\}$
- Org. can address ($s = a$) or dismiss ($s = d$)
- Stakeholders observe decision, decide whether to retaliate
- Retaliate if believe org. made wrong decision...
  - $s = a \rightarrow$ retaliate if $\mu > \theta_a$
  - $s = d \rightarrow$ retaliate if $\mu < \theta_d$
- Confidentiality $\rightarrow$ only info about validity is org. decision

Example

- Organization "stubborn," prefers to dismiss, gains 1
- Wants to avoid retaliation, costs $l \in (0, 1)$

	Retaliate	Don't
Dismiss	1-l	1
Address	$-l$	0

- Stakeholders have limited sway, $l < 1$.
- Rather dismiss w/ retaliation than address $(1 - l > 0)$
- No commitment: only sequentially rational strategy to dismiss

Example

- *Interim payoff*: sender's expected payoff when $s \in S$, common belief μ
- In this example,

$$\widehat{v}(\mu, d) = 1 - l\mathcal{I}(\mu < \theta_d)$$
$$\widehat{v}(\mu, a) = -l\mathcal{I}(\mu > \theta_a).$$

Example

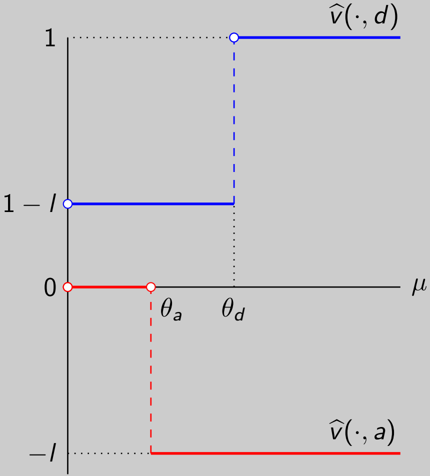

Interim payoff graphs, belief $\mu \equiv \Pr(\omega = f)$

Example

- Sender strategy $\pi(s|\omega)$, for $s \in \{a, d\}$ and $\omega \in \{v, f\}$.
- Receiver beliefs $\{\mu_a, \mu_d\}$, determined by sender's strategy.
- Each $s \in \{a, d\}$ played with some probability
- Probability of μ_s determined by sender strategy, $\tau(\mu_s)$.
- Law of Iterated Expectations ("Bayes-Plausibility")

$$E_\tau[\mu_s] = \mu_0 \iff \tau(\mu_d)\mu_d + \tau(\mu_a)\mu_a = \mu_0$$

- If belief system is BP, some underlying sender strategy induces it
- Can work with beliefs $\{\mu_a, \mu_d\}$ instead of strategy

Example

- Value of belief system $\{\mu_a, \mu_d\}$,

$$\widetilde{v}(\mu_0) = E_\tau[\widehat{v}(\mu_s, s)] = \tau(\mu_d)\widehat{v}(\mu_d, d) + \tau(\mu_a)\widehat{v}(\mu_a, a).$$

- Sender looking for Bayes-Plausible belief system to maximize value
- "Geometric" observation:

$$(\mu_0, \widetilde{v}(\mu_0)) = E_\tau[(\mu_s, \widehat{v}(\mu_s, s))].$$

  If first component averages to prior, second component averages to payoff.
- Can solve the sender problem graphically...

Example

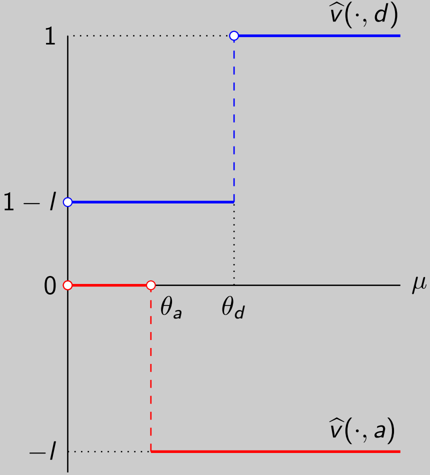

Example

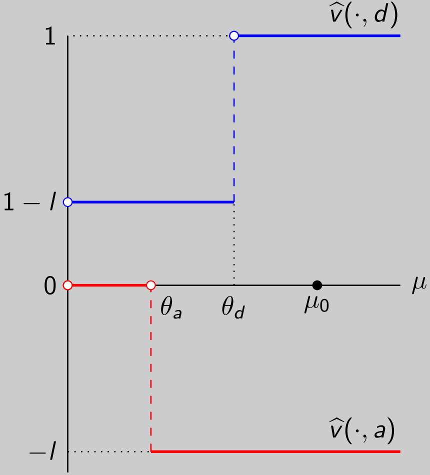

Example

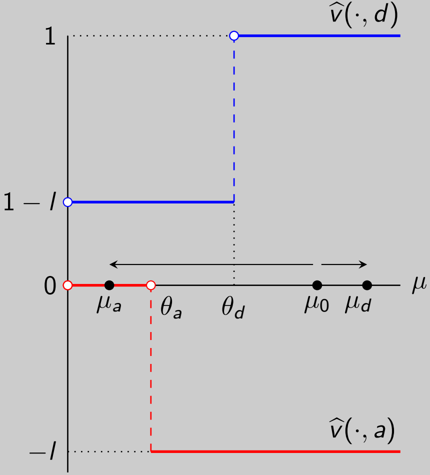

$\widehat{v}(\cdot, d)$
1
$1-l$
0
μ
μ_a
θ_a
θ_d
μ_0
μ_d
$\widehat{v}(\cdot, a)$
$-l$

Example

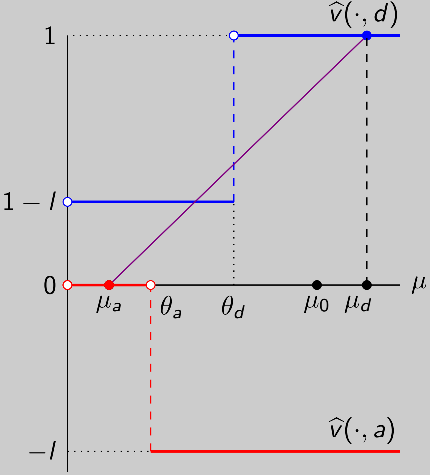
$\widehat{v}(\cdot, d)$
1
$1 - l$
0
$-l$
μ_a
θ_a
θ_d
μ_0
μ_d
μ
$\widehat{v}(\cdot, a)$

Example

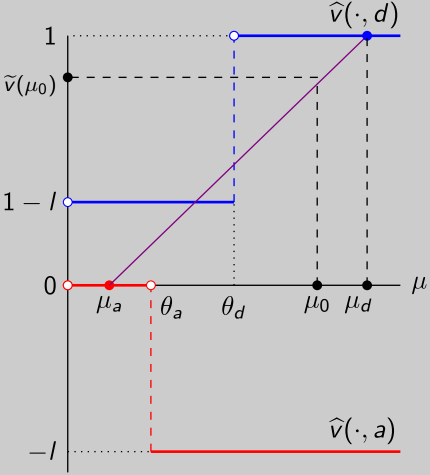

$\widehat{v}(\cdot, d)$
1
$\widetilde{v}(\mu_0)$
$1 - l$
0
μ
μ_a
θ_a
θ_d
μ_0
μ_d
$\widehat{v}(\cdot, a)$
$-l$

Example

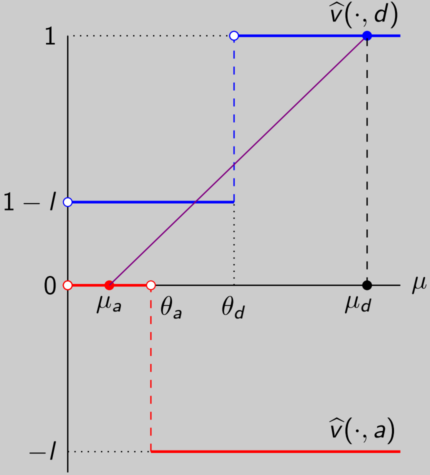

$\widehat{v}(\cdot, d)$
1
$1 - l$
0
μ
μ_a
θ_a
θ_d
μ_d
$\widehat{v}(\cdot, a)$
$-l$

Example

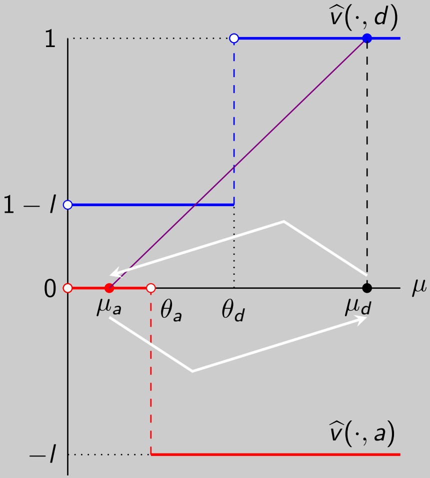

$\widehat{v}(\cdot, d)$
1
$1 - l$
0
μ
μ_a
θ_a
θ_d
μ_d
$\widehat{v}(\cdot, a)$
$-l$

Example

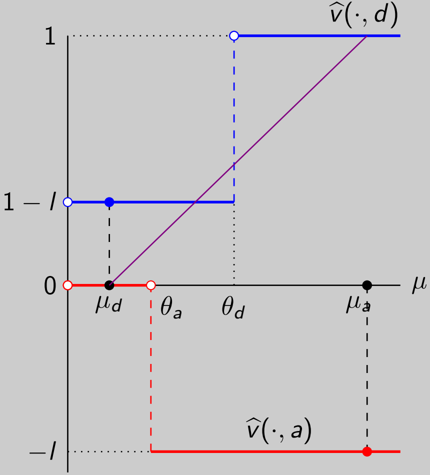

$\widehat{v}(\cdot, d)$
1
$1-l$
0
μ
μ_d
θ_a
θ_d
μ_a
$\widehat{v}(\cdot, a)$
$-l$

Example

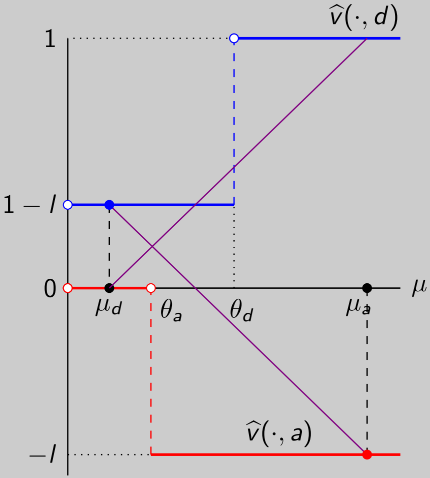

$\widehat{v}(\cdot, d)$
1
$1-l$
0
$-l$
μ
μ_d
θ_a
θ_d
μ_a
$\widehat{v}(\cdot, a)$

Example

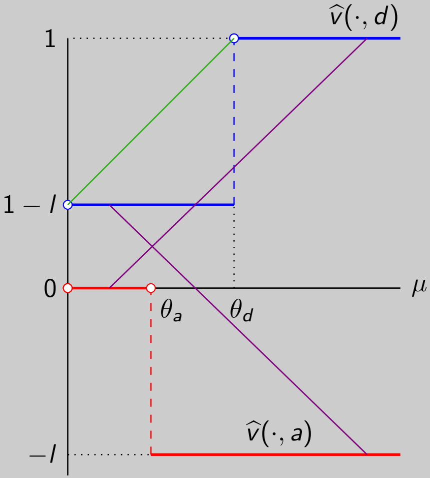

Example

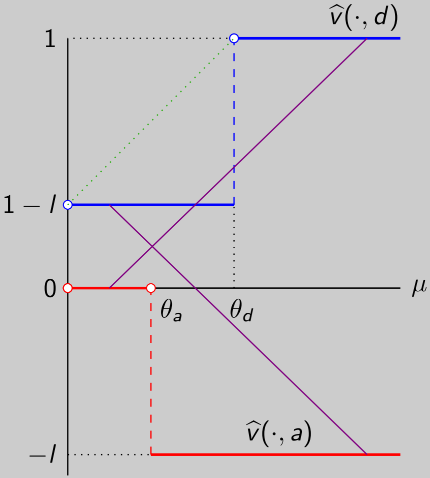

Example

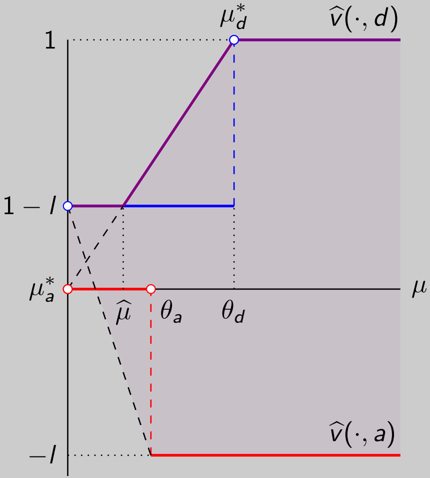

μ_d^*
$\widehat{v}(\cdot, d)$
1
$1 - l$
μ_a^*
μ
$\widehat{\mu}$
θ_a
θ_d
$\widehat{v}(\cdot, a)$
$-l$

Example

Characterization.

- Set of attainable payoffs: payoff graphs + all line segments connecting red graph to blue graph.
- This set called the *topological join* of payoff graphs.
- *Join envelope* is its upper boundary, highest possible payoff at a prior belief
- Beliefs that are joined are optimal belief system
- Join envelope may be convex, generally not concave envelope

Example

Adjudication.

- At moderate beliefs, organization commits to address some valid grievances in order to deter retaliation.
- If prior too high, no retaliation anyway
- If prior too low, must address valid grievances too often, too costly
- At moderate prior, a low probability of (costly) remedy deters retaliation
- Optimal strategy magnifies stakeholders' influence (at moderate priors).
- Must have some influence to begin with ($I > 0$).

Theory

- First part of the theory generalizes the example.
- Results are direct extension.
- Sender interim payoff: $\widehat{v}(\mu, s) = E_\mu[v(\omega, s, \hat{a}(\mu, s))]$
- Graph the sender's payoff functions for each actions $s \in S$, on domain $\Delta(\Omega)$.
- *Result:* Set of attainable sender payoffs is topological join of these graphs, J.
- J is graphs $+$ all line segments connecting *different* graphs
- *Result:* Optimal payoff at each prior is the join envelope,

$$V^{jo}(\mu_0) = \max\{z \,|\, (\mu_0, z) \in J\}.$$

- *Result:* Beliefs that are joined $\rightarrow$ optimal belief system

Theory

Beer-Quiche Example.

- *Tough Sender*: 1 if Beer, 0 if Quiche
- *Wimpy Sender*: 1 if Quiche, 0 if Beer; cost $c \in (0, 1)$ if bullied
- *Receiver*: 0 if leaves alone, $1 - k$ if Bullies wimpy, $-k$ if Bullies tough

Theory

Beer-Quiche Example.

- Only separating eq without commitment, $1 - c > 0$.
- Let $\mu = \Pr(\text{tough})$. Receiver best response is Bully iff $\mu < k$.
- Interim payoffs:

$$\widehat{v}(\mu, Quiche) = (1-\mu)(1 - c\mathcal{I}(\mu < k))$$
$$\widehat{v}(\mu, Beer) = \mu - c(1-\mu)\mathcal{I}(\mu < k)$$

- Interesting case, $k < 1/2$ and $c > k/(1-k)$.

Theory

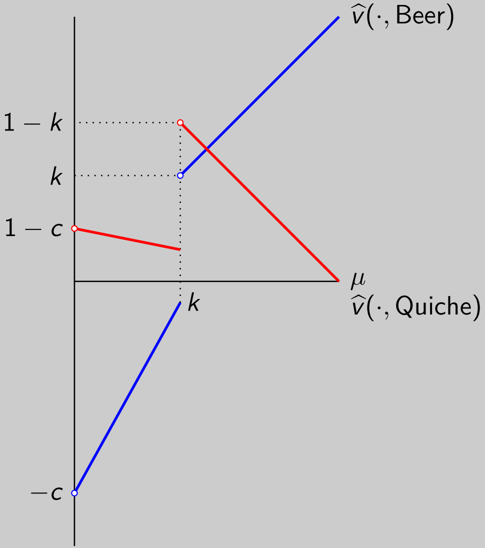

$\widehat{v}(\cdot, \text{Beer})$
$1-k$
k
$1-c$
μ
k
$\widehat{v}(\cdot, \text{Quiche})$
$-c$

Theory

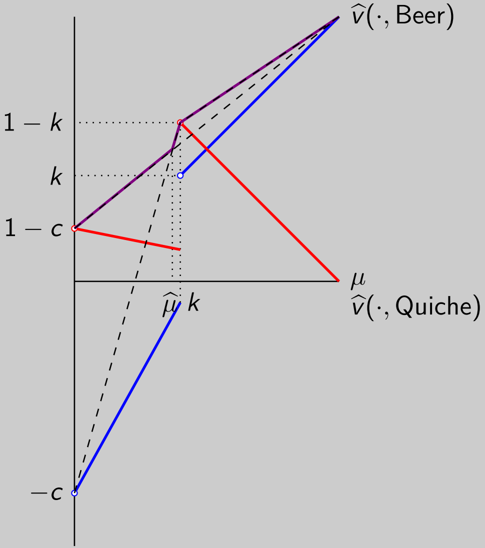
$\widehat{v}(\cdot, \text{Beer})$
$1-k$
k
$1-c$
μ
$\widehat{\mu}$ k
$\widehat{v}(\cdot, \text{Quiche})$
$-c$

Theory

Beer-Quiche Example.

- Separating strategy optimal at low beliefs, $\mu_0 < \widehat{\mu}$.
- For $\mu_0 \in (\widehat{\mu}, k)$, tough quiche and wimpy mixes.
- For $\mu_0 > k$ tough mixes and wimpy brings quiche.
- Partial "reversal" of natural actions.
- Note: for $\mu \leq k$, interim payoff $\widehat{v}(\mu, \text{Quiche}) > \widehat{v}(\mu, \text{Beer})$.

Theory

- Only source of info sender's action
- Maybe prevented from sending extra info (adjudication, trading)
- Maybe sender action *is* a message (grading students, ratings)
- Other settings, sender may have more control over receiver's information
- What would sender do? Is this beneficial for sender (or receiver)?

Theory

- Extend sender's commitment power.
- Along with action, sender commits to communication protocol
- Designs message space M and a joint distribution of public message and action conditional on state $\pi(m, s|\omega)$.
- Transmits information with message, action, or both
- *Result:* Sender cannot increase payoff further by designing info structure. Improvement requires changes in payoff functions, prior belief, or monitoring

Theory

- Observing (m, s), Bayesian update $\mu(m, s)$
- Strategy generates joint distribution of action and belief $\tau(\mu, s)$.
- Decompose joint distribution, $\tau(\mu, s) = \tau_m(\mu)\tau_c(s|\mu)$.
- The belief is Bayes-Plausible, $E_{\tau_m}[\mu] = \mu_0$.
- *Result:* A strategy induces joint $\tau(\mu, s)$ iff marginal τ_m is Bayes-Plausible
- Separability: sender can design a Bayes-Plausible distribution of posteriors $\tau_m(\cdot)$, assign any action to each realized belief $\tau_c(\cdot|\mu)$.

Theory

- Can connect any two points on payoff graphs
- Attain payoff in convex hull of union of graphs
- *Result:* max payoff is concave envelope of graphs
- *Result:* sender commits to select action from highest graph at each belief,

$$\tau_c(s|\mu) > 0 \Rightarrow \widehat{v}(\mu, s) \geq \widehat{v}(\mu, s'), \quad \text{for all} \quad s' \in S.$$

- Such action is *belief-optimal.* Optimal for sender given public belief $\omega \sim \mu$. May not be optimal if sender *knows* the realized state ω.

Example

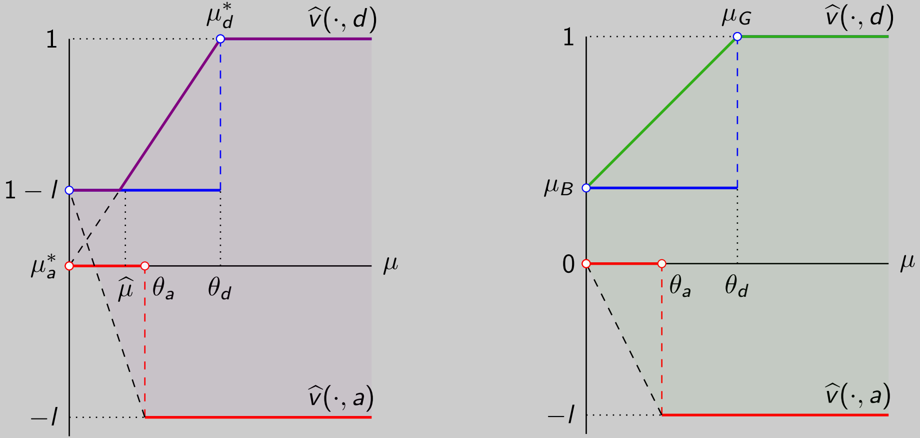

μ_d^*
$\widehat{v}(\cdot, d)$
1
$1-l$
μ_a^*
$\widehat{\mu}$
θ_a
θ_d
μ
$\widehat{v}(\cdot, a)$
$-l$
μ_G
$\widehat{v}(\cdot, d)$
1
μ_B
0
θ_a
θ_d
μ
$\widehat{v}(\cdot, a)$
$-l$

Example

Adjudication.

- In adjudication example, suppose relax confidentiality, allow sender to commit to communication protocol
- Sender always dismisses, uses public message to sometimes deter retaliation
- Undermines the adjudication process (want valid addressed)
- Rationale for confidentiality, beyond privacy concerns

Theory

- To attain "extended commitment" payoff (highest possible), sender can...
  1. Design an experiment that reveals public info before learning state
  2. Delegate action to aligned but *uninformed* intermediary
- (1) allows sender to reveal information according to τ_m
- (2) ensures that at each realized belief, action on highest graph (belief-opt)
- Generally, commitment to communication protocol alone insufficient...

Theory

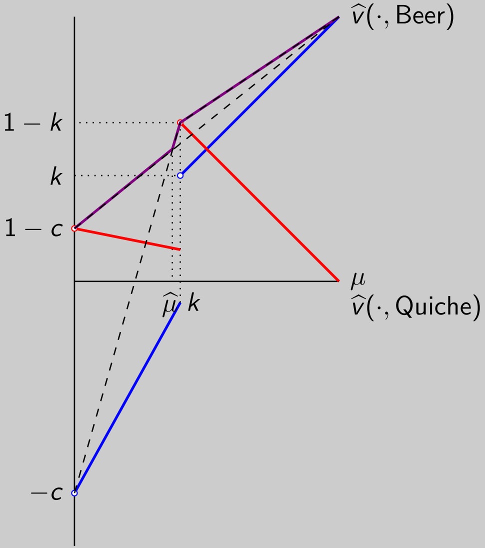

$\widehat{v}(\cdot, \text{Beer})$
$1-k$
k
$1-c$
μ
$\widehat{\mu}$
k
$\widehat{v}(\cdot, \text{Quiche})$
$-c$

Theory

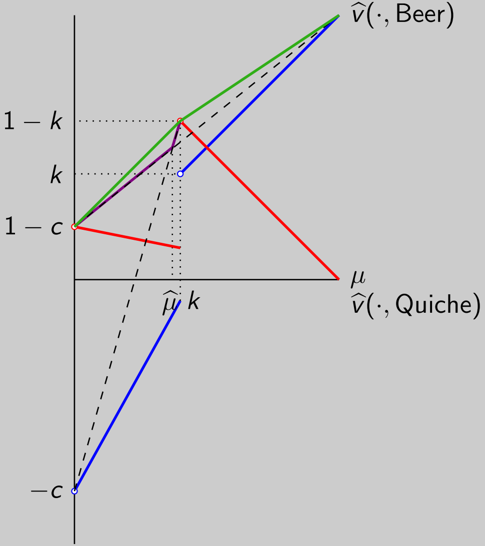

$\widehat{v}(\cdot, \text{Beer})$
$1-k$
k
$1-c$
μ
$\widehat{\mu}$
k
$\widehat{v}(\cdot, \text{Quiche})$
$-c$

Theory

- Concave envelope links $(0, 1-c)$ and $(k, 1-k)$, both on $\widehat{v}(\cdot, \mathit{Quiche})$.
- To achieve concave envelope, sender commits to split belief $\mu_0 \in (0, k)$ into $\{0, k\}$ and select action Quiche.
- Without commitment, tough type always selects Beer.
- Commitment to action required to attain concave envelope, cannot be attained with communication protocol alone.

Theory

- In some environments, concave envelope can be achieved without commitment to communication protocol, only action.
  - Suppose that sender designs a communication protocol, inducing τ_m
  - At each realized posterior μ, sender plays optimal "signaling with commitment" strategy, payoff $V^{jo}(\mu)$.
  - When designing τ_m, expects payoff $V^{jo}(\cdot)$ at each belief.
  - Implies that extended commitment payoff is concave envelope of $V^{jo}(\cdot)$
- *Result:* For all prior beliefs $\mu_0 \in \Delta(\Omega)$, commitment to communication protocol does *not* increase sender payoff if and only if $V^{jo}(\cdot)$ is concave.
- Sufficient condition—all payoff graphs concave.

Application: Financial Rating

Key Feature: nominal rating matters beyond information content

- Designer benefits directly from a favorable rating
  - Financial analysts rewarded more for optimism than accuracy (literature)
  - Professor gives 'A' $\rightarrow$ student happy $\rightarrow$ good evals $\rightarrow$ Dean happy
  - Quality certifier wants to preserve future business
  - Naive receivers, manipulated into favorable responses (finance, hiring)
- Exaggeration costly
  - Financial analysts may be sanctioned for manipulation
  - Professor dislikes inflating grades of undeserving students
  - Risky to certify unsuitable or dangerous product
  - Aversion to self-serving lies, exploitation

Application: Financial Rating

- An asset can be in a good or bad state, $\omega \in \{1, 0\}$, prior $\mu_0 = \Pr(\omega = 1)$
- Analyst chooses a disclosure rule/test $\pi(\cdot|\omega)$, prob rating $s \in \{H, L\}$ given ω
- State of asset observed privately by analyst, rating issued following rule/test
- Continuum of investors, fraction ν naive, $1 - \nu$ sophisticated, dx capital
- Sophisticated observe the disclosure rule and rating, update belief
- Naive believe rating honest, $H \iff \omega = 1$ and $L \iff \omega = 0$

Application: Financial Rating

- Investors draw i.i.d outside options $\theta_i \sim F(\cdot)$, support $[0, 1]$
- Investors decide whether to invest or take outside option.
- Invests if belief that asset is good exceeds outside option (asset value ω)
- Sophisticated invests if *Bayesian update* exceeds outside option, $\mu \geq \theta_i$
- Naive invests if $s = H$
- Aggregate investment at μ is $(1 - \nu)F(\mu) + \nu\mathcal{I}(s = H)$

Application: Financial Rating

- Analyst would like to increase investment in asset $\rightarrow$ use high rating to manipulate naive types
- Analyst would also like to avoid exaggeration.
- If high rating on bad asset, (expected) cost $k > 0$

Application: Financial Rating

Normalizing by $(1-\nu)^{-1}$, interim payoffs

$$\widehat{v}(\mu, H) = F(\mu) + b - c(1-\mu)$$
$$\widehat{v}(\mu, L) = F(\mu)$$

- Concavity/convexity of interim payoffs determined by $F(\cdot)$.
- Position determined by sign$(b-c)$.
- If $b > c$, graph $\widehat{v}(\cdot, H)$ strictly above
- If $b < c$, graphs cross once.

Application: Financial Rating

- Concavity/convexity $\rightarrow$ reveal or conceal from sophisticated investors
  - If $F(\cdot)$ concave, sophisticated investment maximized by concealing
  - If $F(\cdot)$ convex, sophisticated investment maximized by revealing
- Sign$(b - c) \rightarrow$ gain from manipulation (H in bad state).
  - In bad state, high rating increases analyst payoff by $b - c$.
  - If $b < c$, manipulation costly (separation)
  - If $b > c$, manipulation beneficial (pooling on H)
- Sophisticated and naive investors generate incentives that operate through different channels, may reinforce or oppose each other

Application: Financial Rating

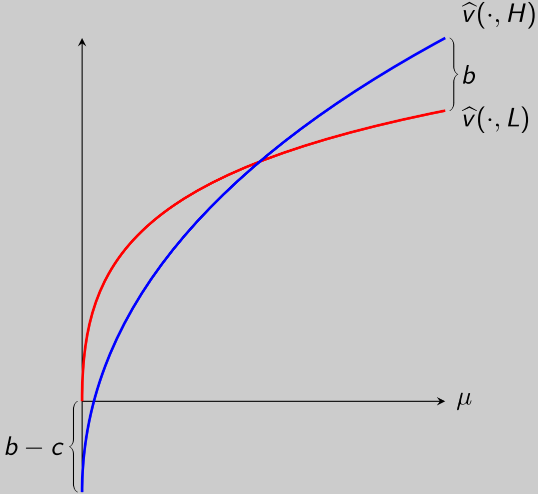

Concave $F(\cdot)$, $b < c$

Application: Financial Rating

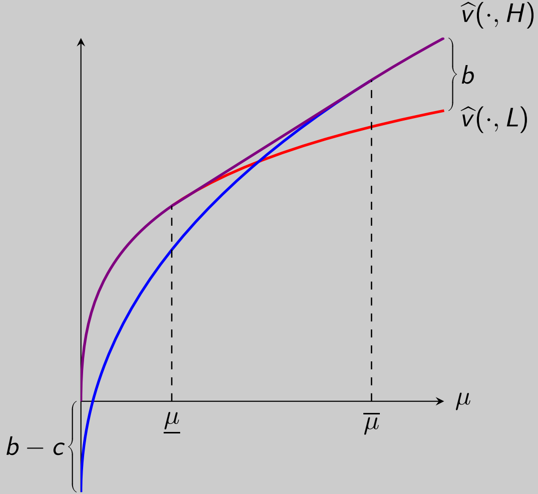

Concave $F(\cdot)$, $b < c$

Application: Financial Rating

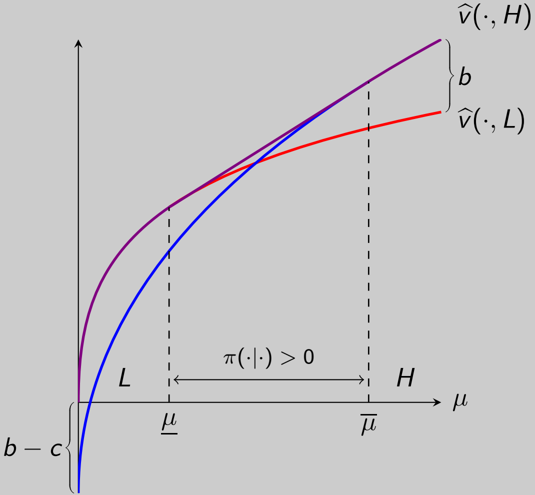

Concave $F(\cdot)$, $b < c$

Application: Financial Rating

- Three forces at play
  1. conceal information from sophisticated
  2. gain b in good state, assign H
  3. avoid cost $b - c < 0$ in bad state, assign L.
- High prior, unlikely pay $b - c \rightarrow$ pool H
- Low prior, unlikely to gain $b \rightarrow$ pool L
- Moderate prior, all three forces$\rightarrow$ some info, not full

Applicaxtion: Financial Rating

- Optimal rating may both overstate $\pi(H|0) > 0$ and understate $\pi(L|1) > 0$
- Incentive to separate comes from fraction of naive investors.
- At a given prior, increase in fraction of naive investors can switch optimal rating from uninformative to informative.

Application: Financial Rating

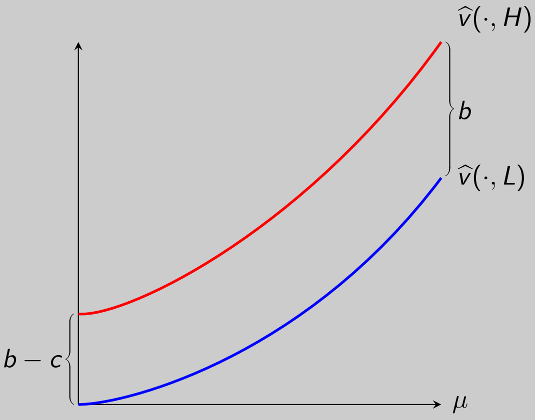

Convex $F(\cdot)$, $b > c$

Application: Financial Rating

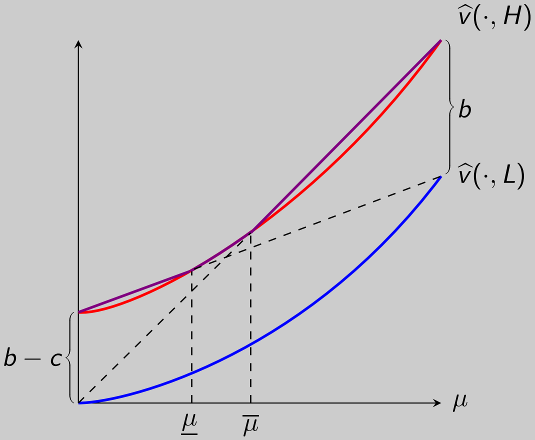

Convex $F(\cdot)$, $b > c$

Application: Financial Rating

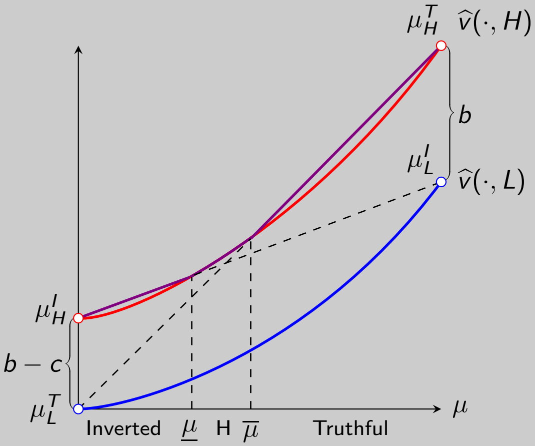

Convex $F(\cdot)$, $b > c$

Application: Financial Rating

- Convexity $\rightarrow$ reveal to max sophisticated investment, separation
- $b > c$ $\rightarrow$ manipulating naive (net) beneficial, pooling on H
- If high fraction of naive $\rightarrow$ pool on H
- Suppose separate, sophisticated investment maximized. Naive invest in state w/ rating H. Assign H to state that is more likely:
  - High prior, H in good state$\rightarrow$truthful
  - Low prior, H in bad state$\rightarrow$inverted

Application: Financial Rating

- Truthful rating more risky than inverted: either both invest or neither.
- Inverted rating, less investment in good state, more in bad.
- Bad state likely, inverted rating better.
- Inversion surprising; examples in literature of highly manipulated ratings; good rating, leads to sophisticated selling and naive buying
- Stronger response than expected without commitment. High rating *discounted*, not inverted.

Application: Financial Rating

Ratings vs. Announcement

- Without commitment, standard taxonomy of equilibria
- In all such equilibria (w/ pooling refined)
  1. Low rating reveals bad state
  2. High rating (weakly) favorable news
  3. Informativeness determined by b vs. c. Shape of $F(\cdot)$ irrelevant.
  4. Higher fraction of naive reduces informativeness
- Optimal rating differs on all four points.

Application: Platform Design

> *Platforms often have a financial incentive to steer customers to particularly profitable products and can use the power of defaults and ordering to accomplish that effectively.*
>
> —Scott-Morton, et al (2019) p. 51

- Platform algorithm orders products, changes incentive to search (steering)
- Platform has superior information about match quality, used by algorithm
- Position also reveals information
- Key idea: how do steering and information provision interact?
- For exposition, present slightly simpler, equivalent model to paper

Application: Platform Design

- Consumer narrowed down a choice to two products, $\{A, B\}$, will buy one
- Product B known, payoff $u \in (0,1)$ (including price).
- Uncertain about A, may be better or worse than B.
- If buys B, payoff $\omega \in \{0,1\}$.
- $\Pr(\omega = 1) = \mu_0$ (throughout μ is belief $\omega = 1$).

Application: Platform Design

- Both products sold on online platform
- To buy or learn about product, consumer must access its "listing."
- Listing for B provides link to purchase product.
- Listing for A provides link, also product information.
- When A's listing accessed, consumer learns true match-quality ω with probability $r < \bar{r}$, otherwise learns nothing ("truth or noise").
- Platform collects commission on sales, $k_A = 1$ and $k_B = 0$.
- Platform paid $f \in (0, \bar{f})$ if sequences A first (for exposition).

Application: Platform Design

- Platform designs algorithm that customizes search results given match ω
- Algorithm affects consumer's ability to learn or buy products ("steering")
- One product listing may be featured at top of page, other buried
- Easy product "positioned" or "sequenced" first, difficult product second
- Because algorithm conditions on match-quality, results reveal information

Application: Platform Design

(0) Platform commits to algorithm, $\pi(s|\omega)$, probability that product $s \in \{A, B\}$ sequenced first, given ω.

(1) Consumer observes first listing . Decides whether to buy first product, or pay search cost $cin(0, \bar{c})$ to access second listing. Buy $\rightarrow$ end

(2) Consumer observes second listing. Decides which product to buy, free recall

Application: Platform Design

- No commitment: A first
- To study commitment, determine probability of selling A when each product sequenced first, assuming consumer's belief is μ upon seeing the first product
- Must solve consumer's (totally standard) search problem. See it

Application: Platform Design

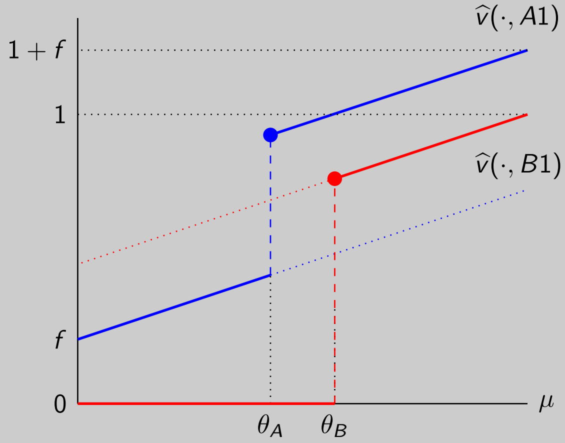

(i) Jumps from change in search strat, (ii) Shift up from fee f, not too big

Application: Platform Design

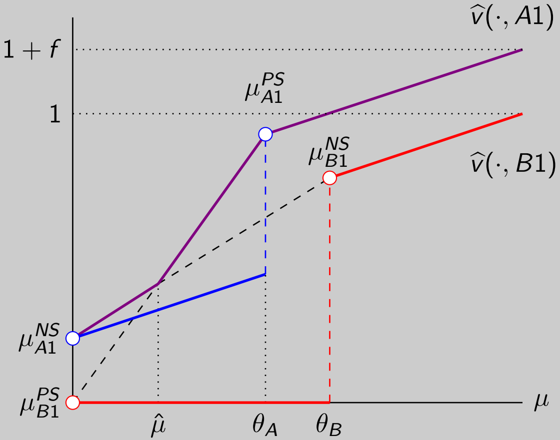

(P): positive sequencing (N): negative sequencing

Application: Platform Design

For $\mu_0 \in (\hat{\mu}, \theta_A)$, optimal algorithm is *positive sequencing*.

- Positive sequencing induces beliefs $\{\mu_{B1} = 0, \mu_{A1} = \theta_A\}$.
- The first position reveals *good news* about the product.
- If B first, consumer knows it matches, buys immediately.
- Uninformed buys A if first, but believes B is better match $(\theta_A < u)$
- Positive sequencing deters search

Application: Platform Design

For $\mu_0 \in (0, \hat{\mu})$, optimal algorithm is *negative sequencing*.

- Negative sequencing induces beliefs $\{\mu_{A1} = 0, \mu_{B1} = \theta_B\}$.
- The first position reveals *bad news* about the product.
- If A first, knows B matches, searches to buy it.
- If B first, consumer searches. Believes A is better match ($\theta_B > u$).
- Here, when consumer searches, believes the second product is likely better.
- Negative sequence encourages search

Application: Platform Design

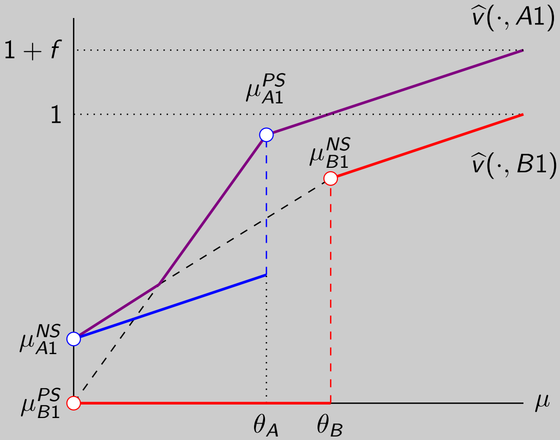

Effect of Search Cost, Small c.

Application: Platform Design

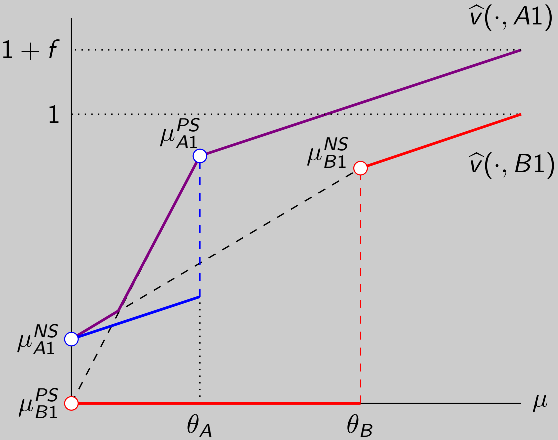

Effect of Search Cost, Large c. θ_A, θ_B spread apart, (PS)+, (NS)-

Application: Platform Design

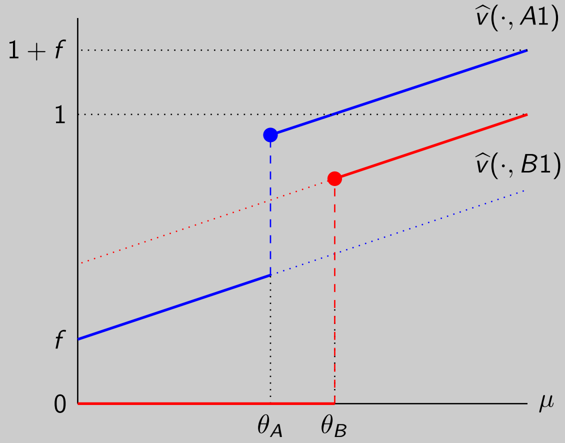

Recommendation System (Extended Commitment)

Application: Platform Design

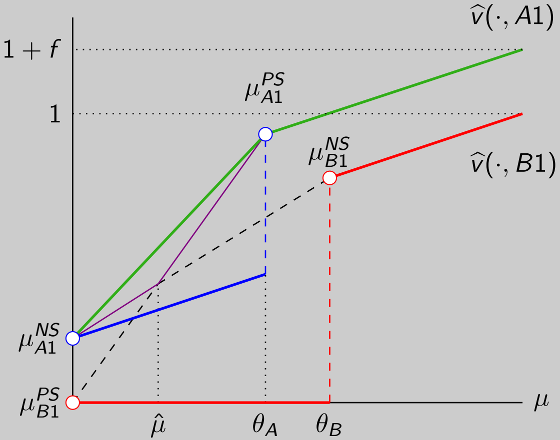

Recommendation System (Extended Commitment)

Application: Platform Design

- Platform benefits from recommendation system.
- Can incentivize most favorable search *and* always place A-first to collect fee
- Optimal recommendation system *hurts* consumer...
- Consumer's ex ante payoff higher with optimal sequencing algorithm
- Gains come from reduction in search cost
- *Result*: PS and NS better for consumer than no commitment ($A1$ always).
  With recommendation system, consumer payoff as if no commitment, worse.
- "Power of defaults and steering" is nuanced

Thanks for listening!

Application: Platform Design

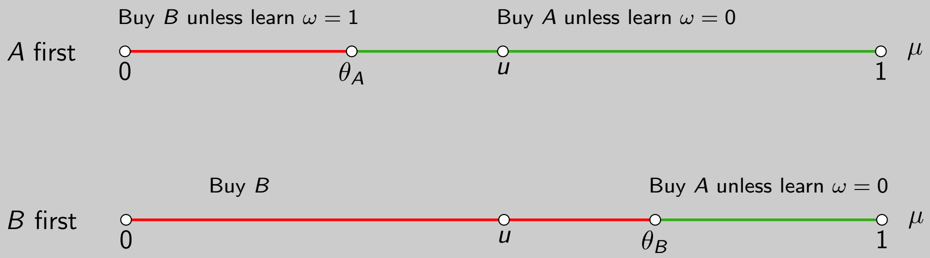

Application: Platform Design

BACK