Too Much of a Good Thing?

The Dynamics of Trust and Loyalty

Johannes Hörner and Anna Sanktjohanser

[1]Yale, TSE (CNRS) and CEPR

[2]Yale, TSE

One World Mathematical Game Theory Seminar

Mariam knows when I need a haircut.
So, she can tell when I am unfaithful.

Mariam knows when I need a haircut.
So, she can tell when I am unfaithful.

And whether that paid off.

Mariam knows when I need a haircut.
So, she can tell when I am unfaithful.

And whether that paid off.

She can adjust the haircut uniformity.


Previous
MECHA-UMA
453 Campbell Ave West Haven, CT (203) 767-0971
Sushi
Cover Photos 1 of 2

Mecha-Uma knows when I need to eat.
So, they can tell when I am unfaithful.

Mecha-Uma knows when I need to eat.
So, they can tell when I am unfaithful.

But **not** whether that paid off.

Mecha-Uma knows when I need to eat.
So, they can tell when I am unfaithful.

But **not** whether that paid off.

They can adjust the meal's pungency.


cannondale
cannondale
cannondale
cannondale

Bob does **not** know when the bike needs to be fixed.
So, he can**not** tell when **I** am unfaithful.

(Nor whether that paid off).

Bob does **not** know when the bike needs to be fixed.
So, he can**not** tell when I am unfaithful.

(Nor whether that paid off).

He can adjust the brakes.

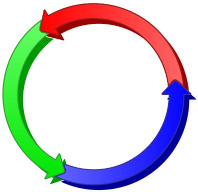

Outside
Option

Buyer: In or Out?

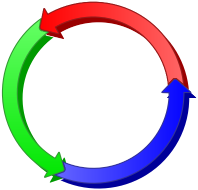

Outside
Option

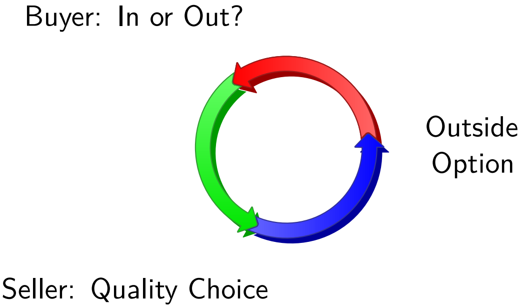
Buyer: In or Out?
Outside Option
Seller: Quality Choice

We focus on the adverse selection vs. moral hazard problem:

When should the buyer come, to discipline the seller?

How should the seller vary quality, to discipline the buyer?

We focus on the adverse selection vs. moral hazard problem:

When should the buyer come, to discipline the seller?

How should the seller vary quality, to discipline the buyer?

We abstract from:

Where does the outside option come from?

How should prices be set?

The Model

Time is discrete, infinite horizon.

One buyer, one seller.

The Model

Time is discrete, infinite horizon.

One buyer, one seller.

Buyer trades in each round, choosing between:

- "*Out*:" An outside option, worth v, i.i.d. : $v \sim F$.
- "*In*:" Going to the seller, worth $q - p$, at seller's cost $c \cdot q$.

The Model

Time is discrete, infinite horizon.

One buyer, one seller.

Buyer trades in each round, choosing between:

"*Out*:" An outside option, worth v, i.i.d. : $v \sim F$.

"*In*:" Going to the seller, worth $q - p$, at seller's cost $c \cdot q$.

Throughout:

The price $p \in (c, 1)$ is paid upfront, and exogenous.

The quality $q \in [0, 1]$ is observable, but noncontractible.

Two Scenarios

The outside option is publicly observed.

The outside option is private information of the buyer.

Three Scenarios

The outside option is publicly observed.

The outside option is private information of the buyer.

The opportunity to trade and, if so, the outside option are private information of the buyer (a few words if time permits).

Observable Outside Option: Histories & Strategies

Public history: $h_t = (v_\tau, a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

v_τ is the value of the outside option in round τ,

a_τ is the buyer's choice, *In* or *Out*,

q_τ is the quality provided by the seller (if applicable).

Observable Outside Option: Histories & Strategies

Public history: $h_t = (v_\tau, a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

v_τ is the value of the outside option in round τ,

a_τ is the buyer's choice, *In* or *Out*,

q_τ is the quality provided by the seller (if applicable).

Buyer's strategy in round t:

$$\sigma_t^B \colon H_t \times \mathbf{R}_+ \to \mathcal{A} := \{In, Out\}.$$

Seller's strategy in round t:

$$\sigma_t^S \colon H_t \times \mathbf{R}_+ \to \mathcal{Q} := [0, 1].$$

Observable Outside Option: Histories & Strategies

Public history: $h_t = (v_\tau, a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

v_τ is the value of the outside option in round τ,

a_τ is the buyer's choice, *In* or *Out*,

q_τ is the quality provided by the seller (if applicable).

Buyer's strategy in round t:

$$\sigma_t^B \colon H_t \times \mathbf{R}_+ \to \mathcal{A} := \{In, Out\}.$$

Seller's strategy in round t:

$$\sigma_t^S \colon H_t \times \mathbf{R}_+ \to \mathcal{Q} := [0, 1].$$

Solution Concept: Subgame-Perfect Equilibrium.

Unobservable Option: Histories & Strategies

Public history: $h_t = (a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

a_τ is the buyer's choice, *In* or *Out*, in round τ, and

q_τ is the quality provided by the seller (if applicable).

Unobservable Option: Histories & Strategies

Public history: $h_t = (a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

a_τ is the buyer's choice, *In* or *Out*, in round τ, and

q_τ is the quality provided by the seller (if applicable).

Private history of the buyer: $h_t^B = h_t \times (v_\tau)_{\tau=1}^{t}$.

Unobservable Option: Histories & Strategies

Public history: $h_t = (a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

a_τ is the buyer's choice, *In* or *Out*, in round τ, and

q_τ is the quality provided by the seller (if applicable).

Private history of the buyer: $h_t^B = h_t \times (v_\tau)_{\tau=1}^{t}$.

Buyer's strategy in round t:

$$\sigma_t^B : H_t^B \to \mathcal{A}.$$

Seller's strategy in round t:

$$\sigma_t^S : H_t \to \mathcal{Q}.$$

Unobservable Option: Histories & Strategies

Public history: $h_t = (a_\tau, q_\tau)_{\tau=1}^{t-1}$, where:

a_τ is the buyer's choice, *In* or *Out*, in round τ, and

q_τ is the quality provided by the seller (if applicable).

Private history of the buyer: $h_t^B = h_t \times (v_\tau)_{\tau=1}^{t}$.

Buyer's strategy in round t:

$$\sigma_t^B \colon H_t^B \to \mathcal{A}.$$

Seller's strategy in round t:

$$\sigma_t^S \colon H_t \to \mathcal{Q}.$$

Solution Concept: Perfect Bayesian Equilibrium.

Benchmarks: First Best & Myopia

To maximize social surplus (buyer + seller payoff):

(i) $q = 1$, and

(ii) the buyer picks *Out* iff

$$v \geq q - cq.$$

Benchmarks: First Best & Myopia

To maximize social surplus (buyer + seller payoff):

(i) $q = 1$, and

(ii) the buyer picks *Out* iff

$$v \geq q - cq.$$

The buyer prefers to pick *Out* iff

$$v \geq q - p.$$

(Myopic behavior). So, she would prefer to come less often.

Maintained Assumptions

If the buyer always acts myopically, the seller prefers to renege on any $q > 0$, **even if doing so ends the relationship**.

(“Myopic behavior isn’t part of an eq’m”).

Maintained Assumptions

If the buyer always acts myopically, the seller prefers to renege on any $q > 0$, **even if doing so ends the relationship**.

("Myopic behavior isn't part of an eq'm").

On average, the outside option exceeds the seller's best offer:

$$\mathbf{E}[v] > 1 - p.$$

("Always *In* isn't part of an eq'm").

Observable Outside Option

Observable Outside Option: Off-Path

The minmax payoff vector $(\mathbf{E}[v], 0)$ is an eq'm payoff vector: autarky.

Wlog, an observable deviation by a player triggers perpetual minmax.

Buyer-Preferred Eq'm: A Candidate

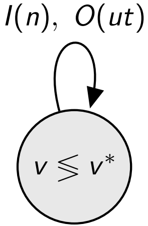

Buyer comes to the seller iff v is below cutoff v^*.

Buyer-Preferred Eq'm: A Candidate

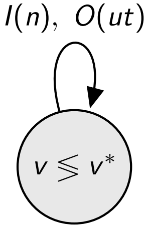

Buyer comes to the seller iff v is below cutoff v^*.

The cutoff v^* is chosen so that the seller IC binds (given q).

The quality $q > 0$ is stationary.

Can we do Better?

At v^*, the buyer would prefer *Out*: $v^* > q - p$.

Yet, she comes because she is the seller's obligee.

Can we do Better?

At v^*, the buyer would prefer *Out*: $v^* > q - p$.

Yet, she comes because she is the seller's obligee.

No such reciprocation is due at the first visit.

Observable Outside Option: A Better Eq'm

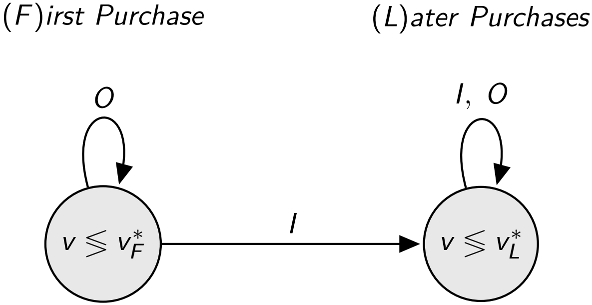

Observable Outside Option: A Better Eq'm

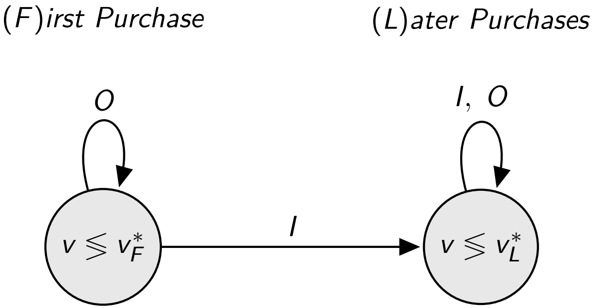

At v_F^*, buyer indifferent between *In* and *Out*, given the transitions:

$$v_F^* < q - p$$

Observable Outside Option: A Better Eq'm

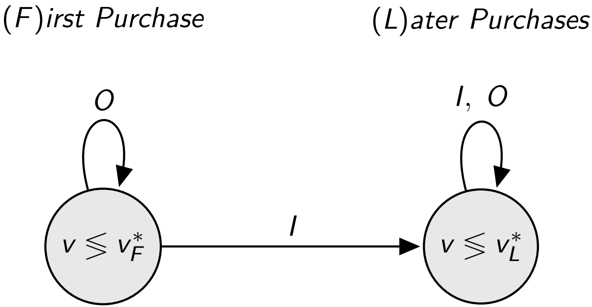

At v_F^*, buyer indifferent between *In* and *Out*, given the transitions:

$$v_F^* < q - p < v_L^*.$$

Observable Outside Option: A Better Eq'm

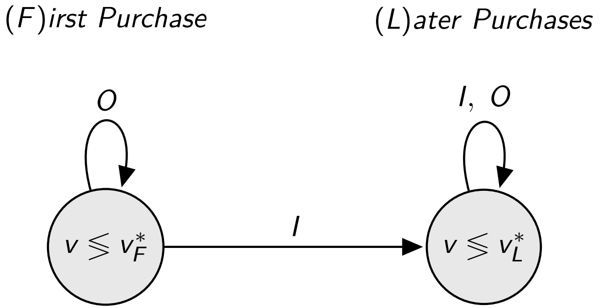

At v_F^*, buyer indifferent between *In* and *Out*, given the transitions:

$$v_F^* < q - p < v_L^*.$$

Seller IC binds after initial purchase.

Seller's quality q is constant.

Idea of Proof: Dynamic Programming On Path

Promised seller's payoff S as a state variable.

Idea of Proof: Dynamic Programming On Path

Promised seller's payoff S as a state variable.

$$\begin{aligned}\mathcal{B}\colon \mathbf{R}_+ &\to \mathbf{R}_+ \cup \{-\infty\}\\ S &\mapsto \sup\left\{\text{Buyer's Reward} + \text{Continuation Payoff } \mathcal{B}(S')\right\}\end{aligned}$$

over *In* or *Out*, q and S', as a function of (S, v), s.t.

Quality q is incentive compatible (Seller IC),

The seller's payoff is at least S (Seller PK).

Idea of Proof: Dynamic Programming On Path

Promised seller's payoff S as a state variable.

$$\begin{aligned}\mathcal{B}\colon \mathbf{R}_+ &\to \mathbf{R}_+ \cup \{-\infty\} \\ S &\mapsto \sup \{\text{Buyer's Reward} + \text{Continuation Payoff } \mathcal{B}(S')\}\end{aligned}$$

over v^* s.t. *In* iff $v \leq v^*$, $q \in \mathcal{Q}$ and $S' \in \{S^I, S^O\}$, s.t.

Quality q is incentive compatible (Seller IC),

The seller's payoff is at least S (Seller PK).

Idea of Proof: Dynamic Programming On Path

Explicitly:

$$\mathcal{B}\colon \mathbf{R}_+ \to \mathbf{R}_+ \cup \{-\infty\}$$
$$S \mapsto \sup \Big\{(1-\delta)\,(F(v^*)(q-p) + (1-F(v^*))\mathbf{E}[v|v \geq v^*]) + \delta\left(F(v^*)\mathcal{B}(S^I) + (1-F(v^*))\mathcal{B}(S^O)\right)\Big\}$$

over v^*, q, S^I, S^O, as a function of S, s.t.

Quality q is incentive compatible (Seller IC),

The seller's payoff is at least S (Seller PK).

Observable Outside Option: Eq'm Pareto Frontier

The buyer-preferred eq'm is a special case of the following.

Proposition (Eq'm Pareto frontier when $q = 1$)

Fix a Pareto-optimal eq'm payoff vector (B, S). *If:*

$S \geq \tilde{S}$: $S^I = S^O = S$ ("one state automaton")

$S < \tilde{S}$: $S^O = S$ *and* $S^I = \tilde{S}$ ("two state automaton")

where $\tilde{S} := (1 - \delta)c/\delta$.

Observable Outside Option: Eq'm Pareto Frontier

The buyer-preferred eq'm is a special case of the following.

Proposition (Eq'm Pareto frontier when $q = 1$)

Fix a Pareto-optimal eq'm payoff vector (B, S). *If:*

$S \geq \tilde{S}$: $S^I = S^O = S$ ("one state automaton")

$S < \tilde{S}$: $S^O = S$ *and* $S^I = \tilde{S}$ ("two state automaton")

where $\tilde{S} := (1 - \delta)c/\delta$.

When not yet locked in (low S), the buyer only delivers higher rents to the seller once she has come.

Observable Outside Option: Eq'm Pareto Frontier

The buyer-preferred eq'm is a special case of the following.

Proposition (Eq'm Pareto frontier when $q = 1$)

Fix a Pareto-optimal eq'm payoff vector (B, S). *If:*

$S \geq \tilde{S}$: $S^I = S^O = S$ ("one state automaton")

$S < \tilde{S}$: $S^O = S$ *and* $S^I = \tilde{S}$ ("two state automaton")

where $\tilde{S} := (1 - \delta)c/\delta$.

When not yet locked in (low S), the buyer only delivers higher rents to the seller once she has come.

Quality need not be one in general, even along the frontier.

Observable Outside Option: Eq'm Pareto Frontier

Proposition

Fix a Pareto-optimal eq'm payoff vector (B, S). *If:*

$S \geq \tilde{S}$: $S^I = S^O = S$ ("one state automaton")

$S < \tilde{S}$: $S^O = S$ *and* $S^I = \tilde{S}$ ("two state automaton")

where

$$\tilde{S} := (1-\delta)c\tilde{q}/\delta,$$

and

$$\tilde{q} := \min\left\{\frac{\delta F(p(1-c)/c)p}{c(1-\delta)+c\delta F(p(1-c)/c)}, 1\right\}.$$

Observable Outside Option When $S \geq \tilde{S}$

One State

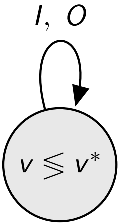

Here,

$$v^* > q - p.$$

Observable Outside Option When $S \geq \tilde{S}$

One State

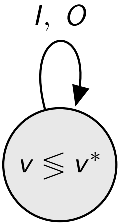

Here,

$$v^* > q - p.$$

Seller IC does not bind: PK and optimality pin down q and v^*.

Observable Outside Option When $S \geq \tilde{S}$

One State

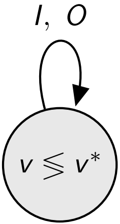

Here,

$$v^* > q - p.$$

Seller IC does not bind: PK and optimality pin down q and v^*.

Seller's quality q is constant.

Observable Outside Option When $S < \tilde{S}$

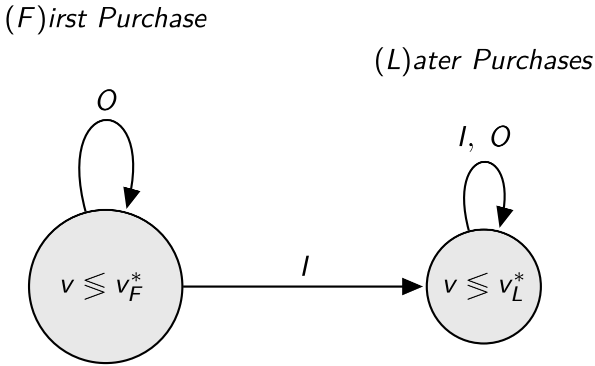

$v_F^* < v_L^*$ and $v_L^* > q - p$.

Observable Outside Option When $S < \tilde{S}$

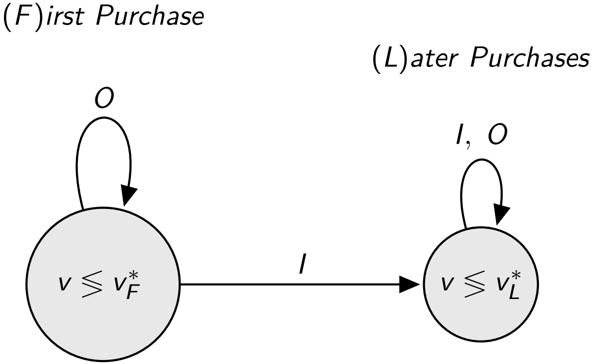

$v_F^* < v_L^*$ and $v_L^* > q - p$.

Seller IC binds after initial purchase.

Observable Outside Option When $S < \tilde{S}$

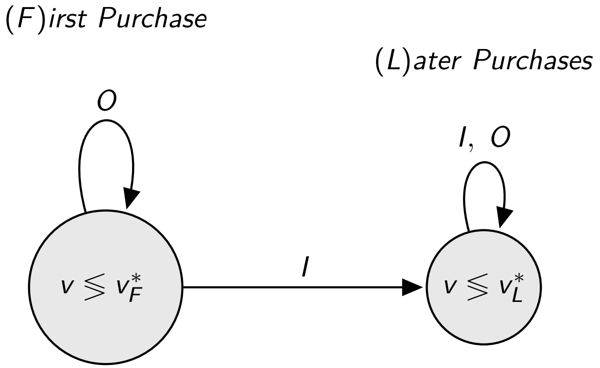

$v_F^* < v_L^*$ and $v_L^* > q - p$.

Seller IC binds after initial purchase.

Seller's quality q is constant.

Observable Outside Option: Summary

Stationary behavior after the first visit.

The buyer:

- waits “too long” before the first purchase, and
- comes “more often than she would like to” thereafter.

Unobservable Outside Option

Unobservable Outside Option: Off Path

Perpetual minmaxing is still an equilibrium.

Unobservable Outside Option: Off Path

Perpetual minmaxing is still an equilibrium.

But only the seller's deviations are observed, and so trigger autarky.

Unobservable Outside Option: Off Path

Perpetual minmaxing is still an equilibrium.

But only the seller's deviations are observed, and so trigger autarky.

The buyer needs to be motivated **on path**: Stationarity is lost.

Buyer-Preferred Two-State Automaton

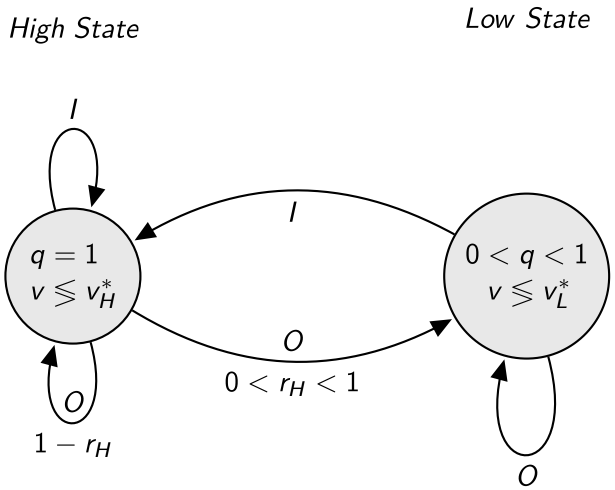

where $1 - p < v_H^*$ and $v_L^* < v_H^*$.

Takeaways

Punishment is moderate: $q > 0$.

The seller is "slow to anger, quick to forgive."

How to Improve?

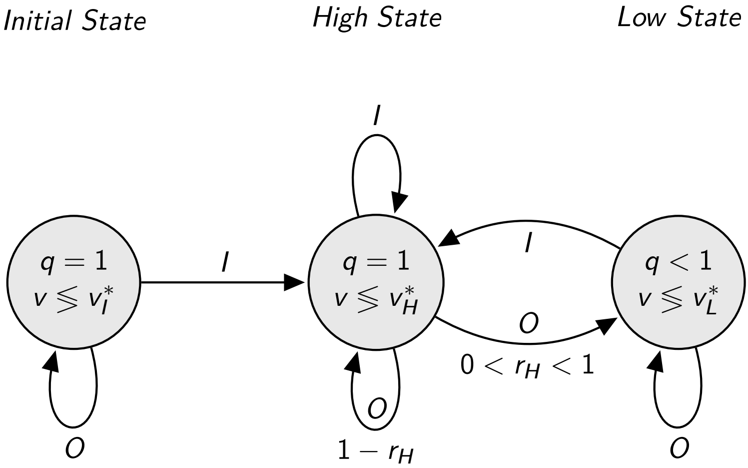

where $v_I^* < 1 - p < v_H^*$ and and $v_L^* < v_H^*$.

Takeaways

Transient consideration stage, before the loyalty loop starts.

How to Improve?

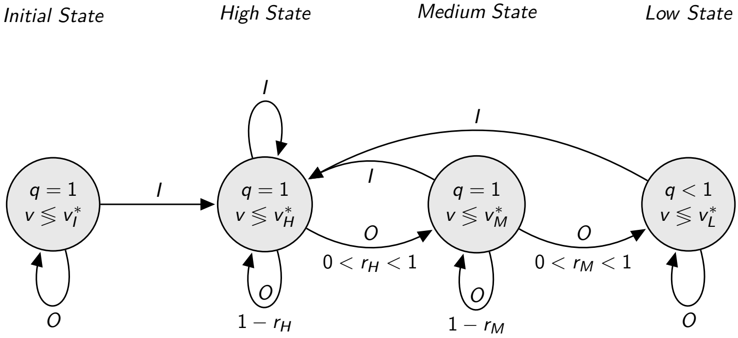

where $v_I^* < 1 - p < v_H^* < v_M^*$ and and $v_L^* < v_H^*$.

Takeaways

Optimality requires calibrated strategies (not a 2-state automaton).

Equilibrium Payoff Set

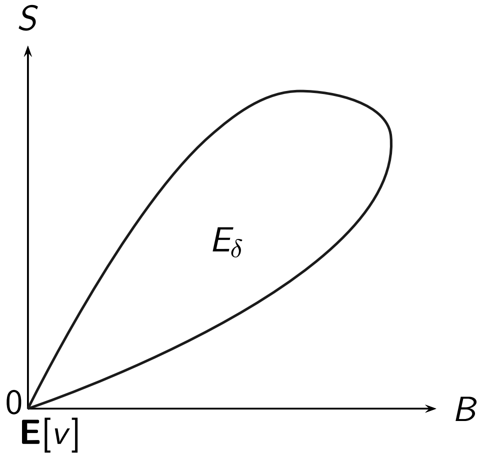

Equilibrium Payoff Set

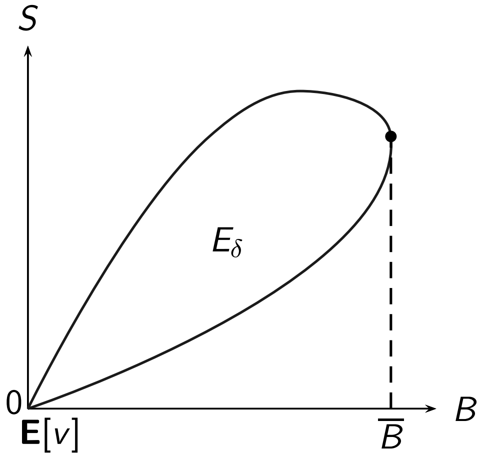

Buyer-Preferred Eq'm –Outside

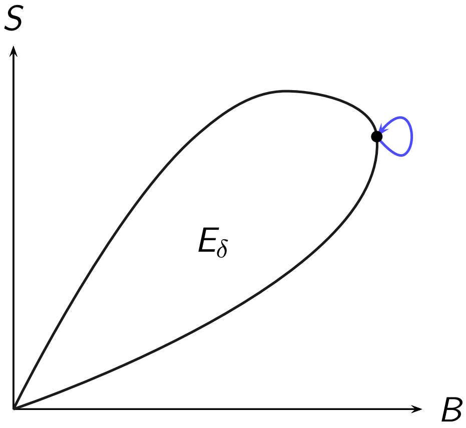

Buyer-Preferred Eq'm –Inside

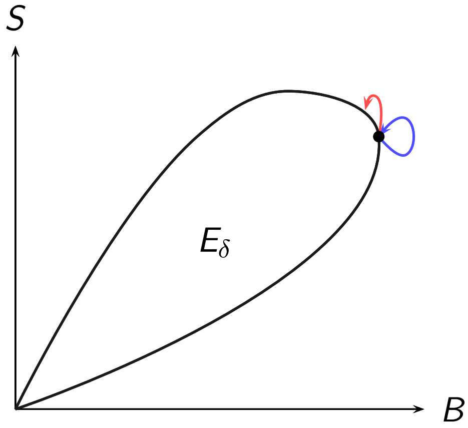

Buyer-Preferred Eq'm –Outside

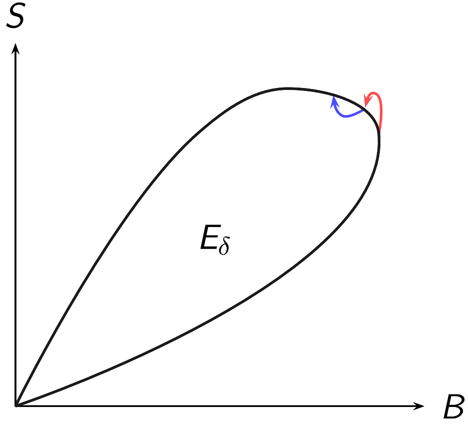

Buyer-Preferred Eq'm –Outside

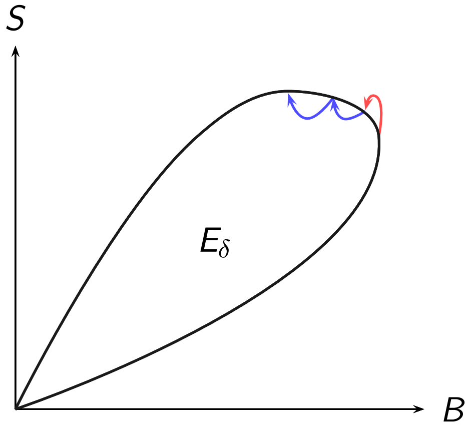

Buyer-Preferred Eq'm –Outside

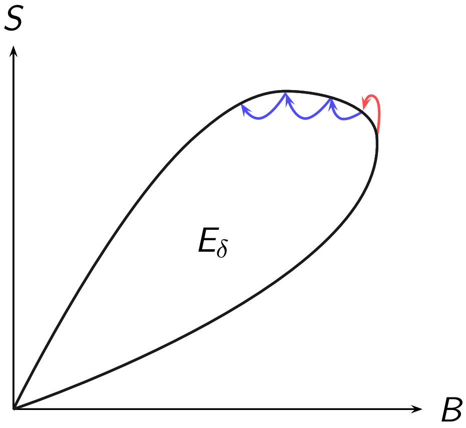

Buyer-Preferred Eq'm –Outside

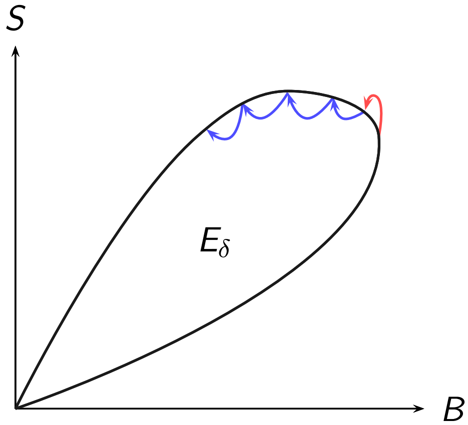

Buyer-Preferred Eq'm –Outside

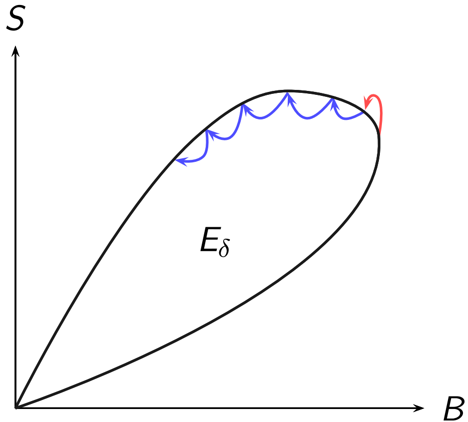

Buyer-Preferred Eq'm –Outside

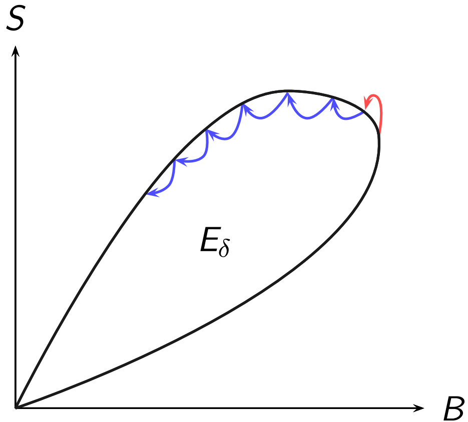

Buyer-Preferred Eq'm –Inside

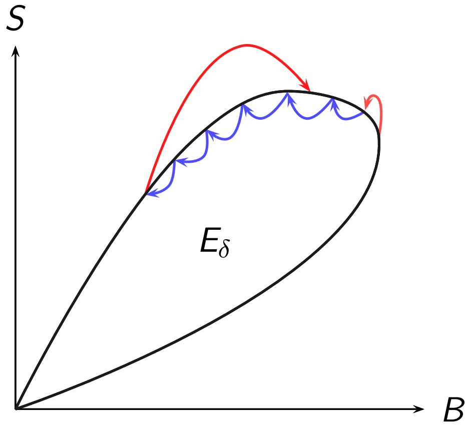

Buyer-Preferred Eq'm –Inside

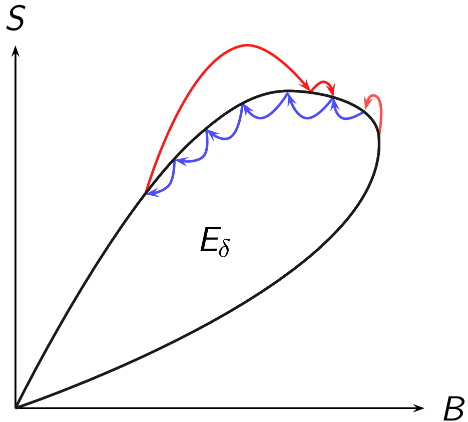

Buyer-Preferred Eq'm –Inside

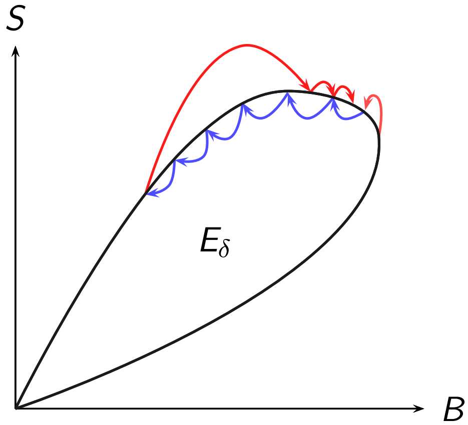

Buyer-Preferred Eq'm –Inside

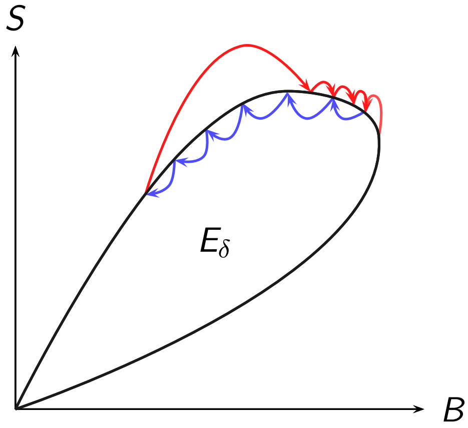

Buyer Preferred Eq'm: Transitions

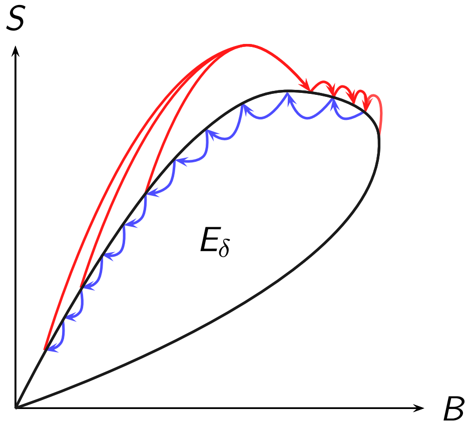

Buyer Preferred Eq'm: Transitions

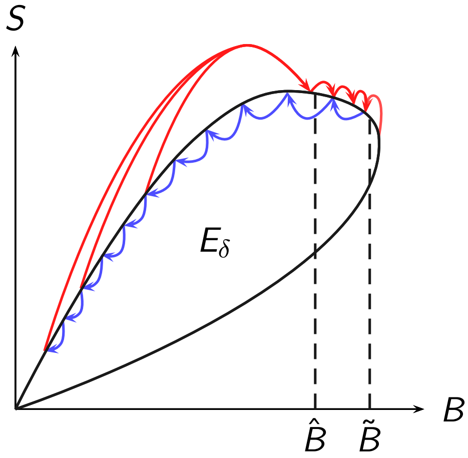

Unobservable Outside Option: Continuation Payoffs

Proposition

Consider the upper boundary of E_δ.

- $B^I(B)$ *and* $B^O(B)$ *are increasing in* B*, with* $B^O(B) \leq B$.
- *Their ranges are* $[\hat{B}, \tilde{B}]$ *and* $[\mathbf{E}[v], \overline{B}]$*, respectively.*

Unobservable Outside Option: Continuation Payoffs

Proposition

Consider the upper boundary of E_δ.

- $B^I(B)$ *and* $B^O(B)$ *are increasing in* B*, with* $B^O(B) \leq B$*.*
- *Their ranges are* $[\hat{B}, \tilde{B}]$ *and* $[\mathbf{E}[v], \overline{B}]$*, respectively.*

That is, payoffs in the range $(\tilde{B}, \overline{B}]$ are transient.

Unobservable Outside Option: Constraints

Proposition

Consider the upper boundary of E_δ.

There exist cutoffs B^q, B^{IC}, with $B^q < B^{IC}$ and $B^{IC} \in (\hat{B}, \tilde{B})$ s.t.

1. *The seller IC binds iff $B_t \geq B^{IC}$,*
2. *$q = 1$ iff $B_t \geq B^q$.*

Further, $B_t < B^q$ i.o., and $B^{IC} < B_t$ i.o.

Upper Boundary of E_δ: Constraints

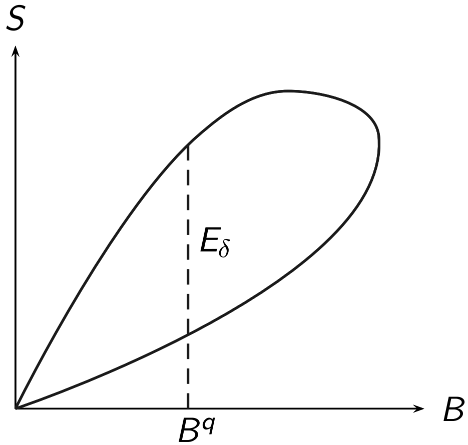

Upper Boundary of E_δ: Constraints

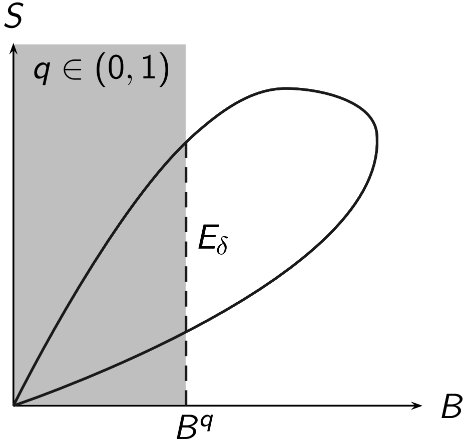

Upper Boundary of E_δ: Constraints

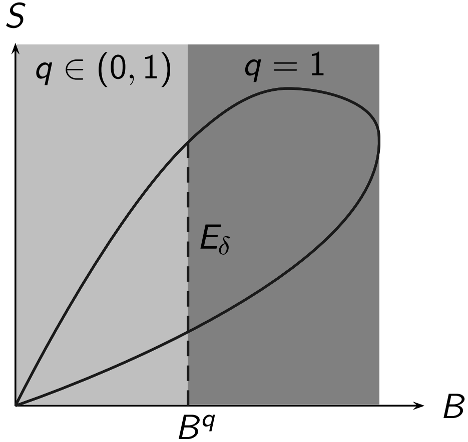

Upper Boundary of E_δ: Constraints

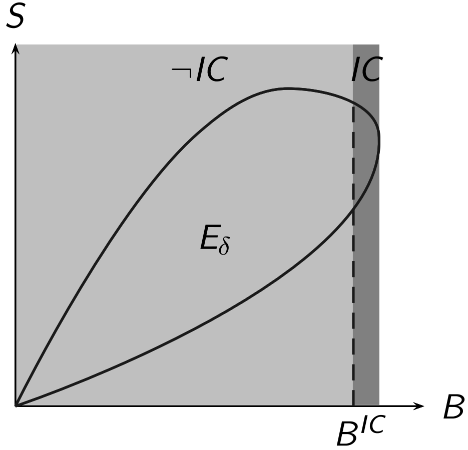

Upper Boundary of E_δ: Constraints

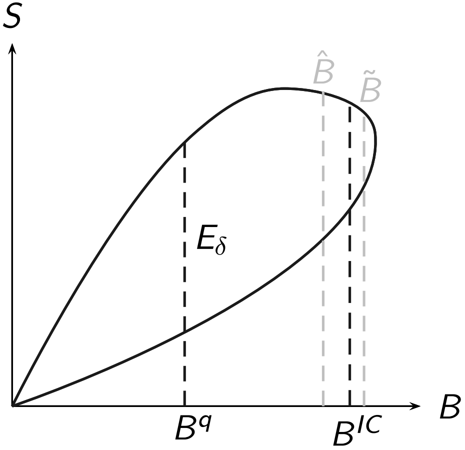

Unobservable Outside Option: Cutoff and Quality

Proposition

Consider the upper boundary of E_δ.

- *The seller's quality $q(B)$ is increasing in B.*
- *The buyer's cutoff $v^*(B)$ is single-peaked in B.*

Quality hits its maximum ($= 1$) when the cutoff hits its peak.

Upper Boundary of E_δ: Cutoff and Quality

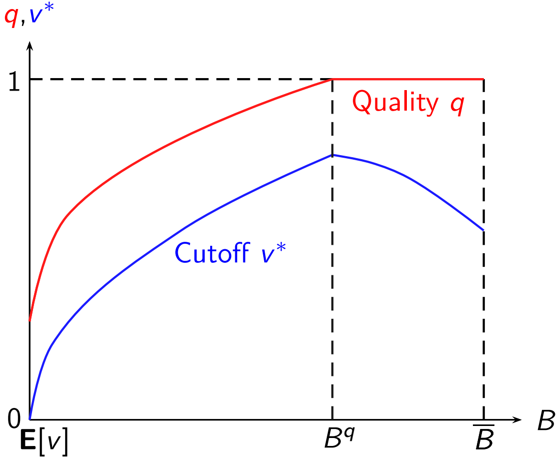

Unobservable Outside Option: Idea of Proof

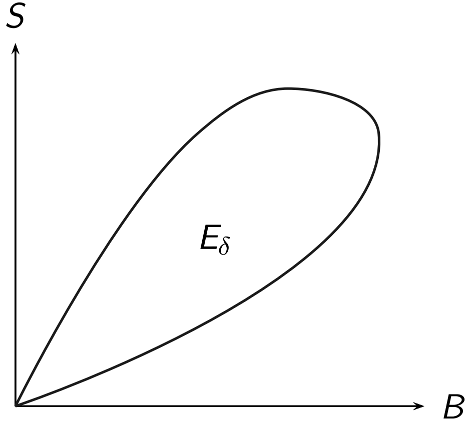

Unobservable Outside Option: Idea of Proof

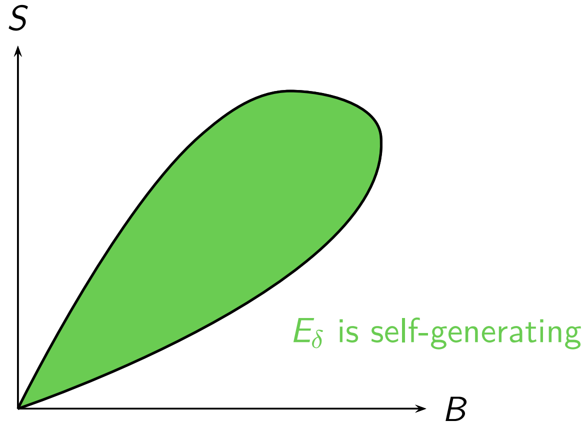

Unobservable Outside Option: Idea of Proof

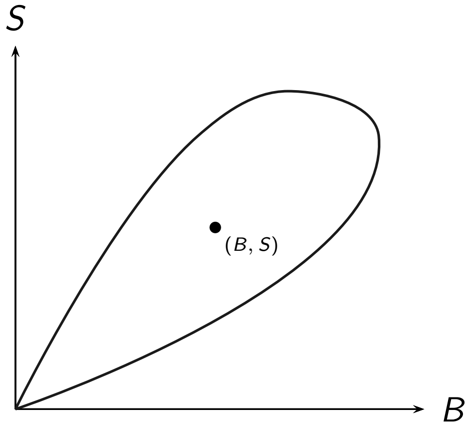

Unobservable Outside Option: Idea of Proof

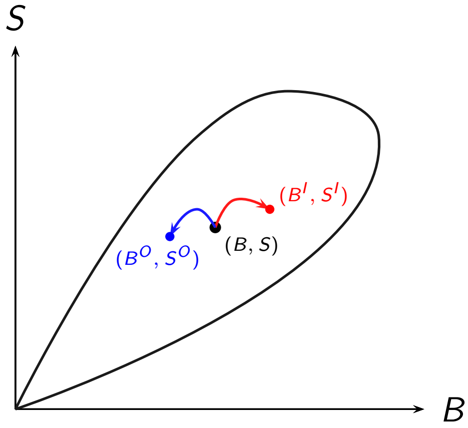

Unobservable Outside Option: Idea of Proof

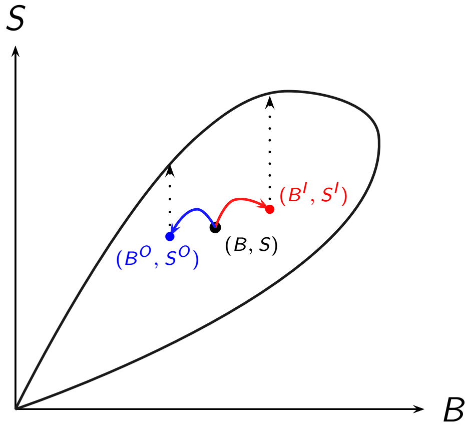

Unobservable Outside Option: Idea of Proof

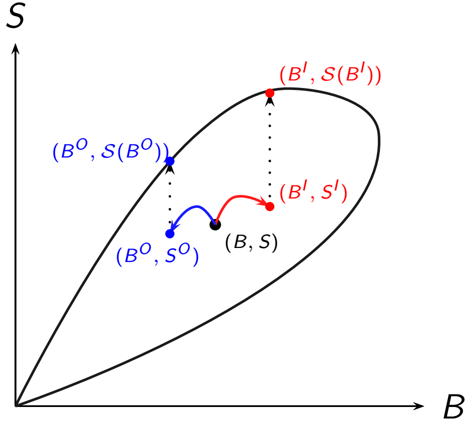

Unobservable Outside Option: Idea of Proof

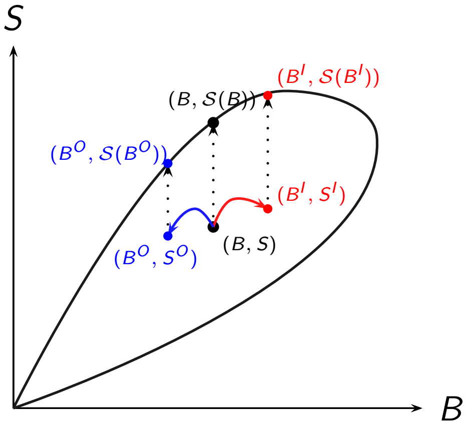

Unobservable Outside Option: Idea of Proof

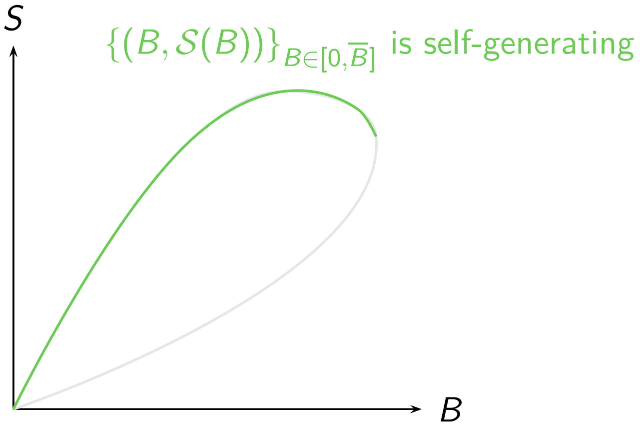

Idea of Proof: Dynamic Programming...

... with one additional constraint. Define:

$$\mathcal{S} \colon \mathbf{R}_+ \to \mathbf{R}_+ \cup \{-\infty\}$$
$$B \mapsto \sup \{\text{Seller's Reward} + \text{Continuation Payoff } \mathcal{S}(\cdot)\}$$

over v^*, q, B^I, B^O, as a function of B, s.t.

Seller IC,

Buyer PK, and

v^* is s.t. $(1-\delta)v^* + \delta B^O = (1-\delta)(q-p) + \delta B^I$.

Idea of Proof: Dynamic Programming...

... with one additional constraint. Define:

$$
\begin{aligned}
\mathcal{S}\colon \mathbf{R}_+ &\to \mathbf{R}_+ \cup \{-\infty\} \\
B &\mapsto \sup \Big\{ (1-\delta)\, F(v^*)(p - cq) \\
&\qquad + \delta \left(F(v^*)\mathcal{S}(B^I) + (1 - F(v^*))\mathcal{S}(B^O) \right) \Big\}
\end{aligned}
$$

over v^*, q, B^I, B^O, as a function of B, s.t.

Seller IC,

Buyer PK, and

v^* is s.t. $(1-\delta)v^* + \delta B^O = (1-\delta)(q-p) + \delta B^I$.

Unobservable Outside Option: Asymptotics

Out begets *Out*:

$$\lim_{B \to \mathbf{E}[v]} v^*(B) = 0.$$

Unobservable Outside Option: Asymptotics

Out begets *Out*:

$$\lim_{B \to \mathbf{E}[v]} v^*(B) = 0.$$

But a single *In* leads to a jump:

$$B > 0 \Rightarrow B^I(B) \geq \hat{B}.$$

Unobservable Outside Option: Asymptotics

Out begets *Out*:

$$\lim_{B \to \mathbf{E}[v]} v^*(B) = 0.$$

But a single *In* leads to a jump:

$$B > 0 \Rightarrow B^I(B) \geq \hat{B}.$$

Incremental, likely penalty vs. radical, unlikely reward.

Unobservable Outside Option: Asymptotics

Out begets *Out*:

$$\lim_{B \to \mathbf{E}[v]} v^*(B) = 0.$$

But a single *In* leads to a jump:

$$B > 0 \Rightarrow B^I(B) \geq \hat{B}.$$

Incremental, likely penalty vs. radical, unlikely reward.

Theorem

It holds that:

$$\underline{\lim}_{n \to \infty} B_t = \mathbf{E}[v] \text{ and } \overline{\lim}_{n \to \infty} B_t = \tilde{B}.$$

Observable Outside Option

Buyer waits "too long" before first purchase.

Buyer comes more often than she would like to after first purchase.

Stationary after first purchase, hence trivial asymptotics.

Unobservable Outside Option

Buyer waits "too long" before first purchase.

Buyer comes more often than she would like to after first purchase.

"Slow to anger, quick to forgive:" Recurrent between near autarky and Pareto frontier.

“Nothing is Observed”

Need to Trade and Value to Trade both Private Information

Similar to the previous case:

As if some outside options are “too good to refuse.”

“Nothing is Observed”

Need to Trade and Value to Trade both Private Information

Similar to the previous case:

As if some outside options are “too good to refuse.”

Buyer’s best payoff is non-monotonic in likelihood of this event.

Buyer-Preferred Eq'm with Increasing Atom

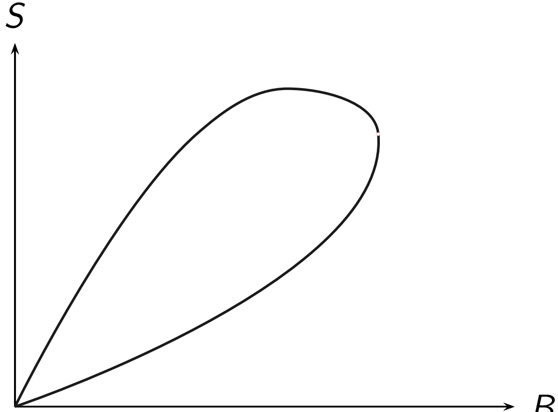

Buyer-Preferred Eq'm with Increasing Atom

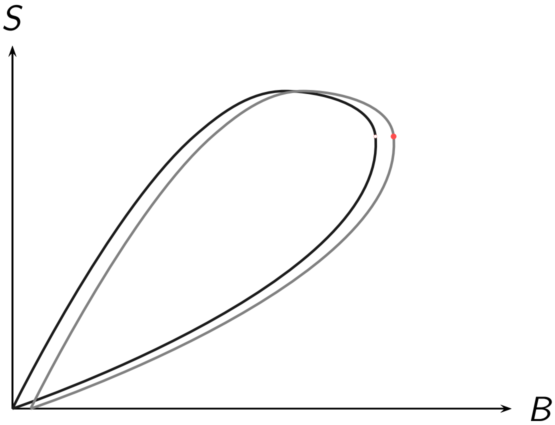

Buyer-Preferred Eq'm with Increasing Atom

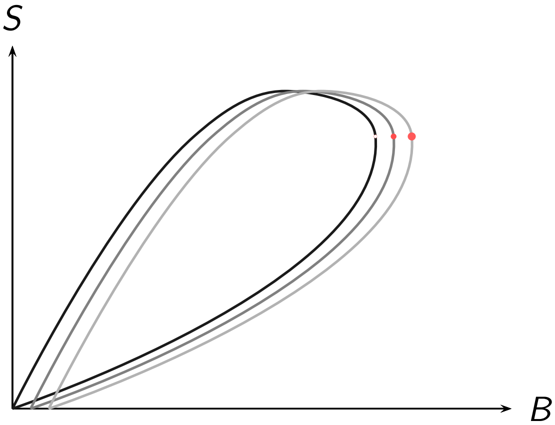

Buyer-Preferred Eq'm with Increasing Atom

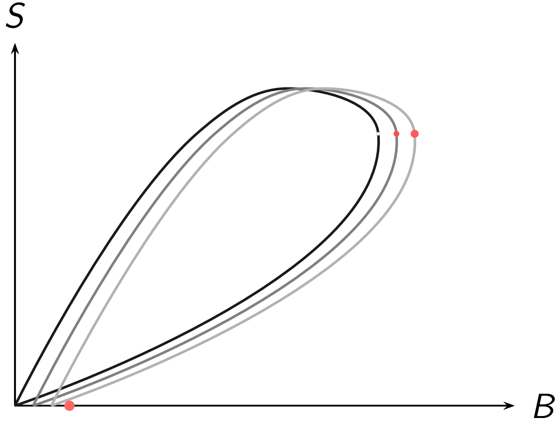

Buyer-Preferred Eq'm with Increasing Atom

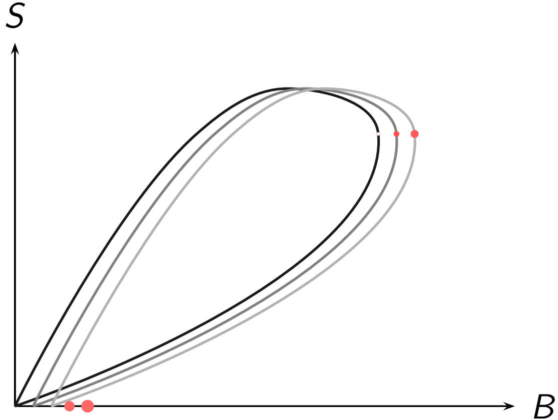

Buyer-Preferred Eq'm with Increasing Atom

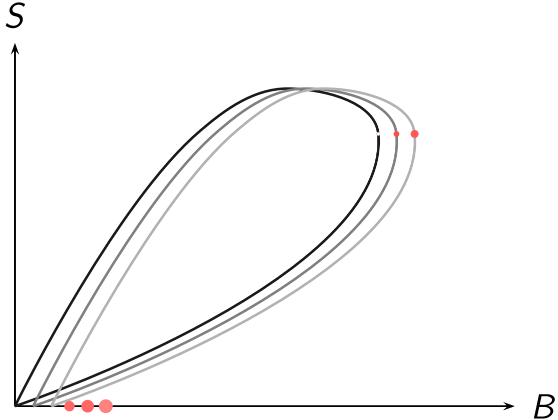

Buyer-Preferred Eq'm with Increasing Atom

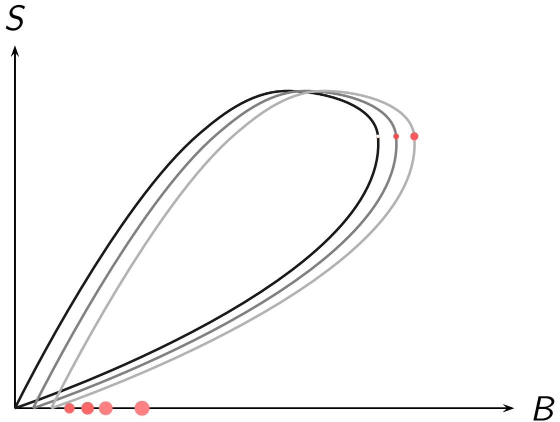

Buyer-Preferred Eq'm with Increasing Atom

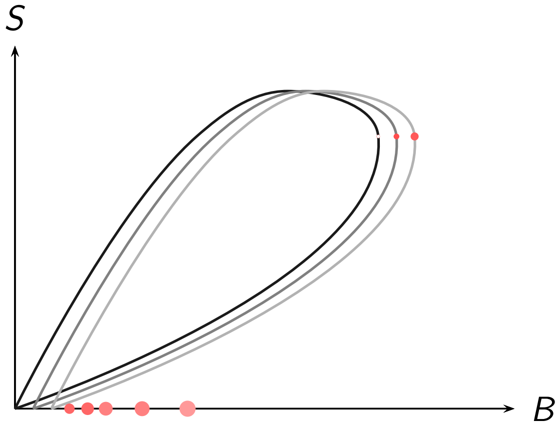

Buyer-Preferred Eq'm with Increasing Atom

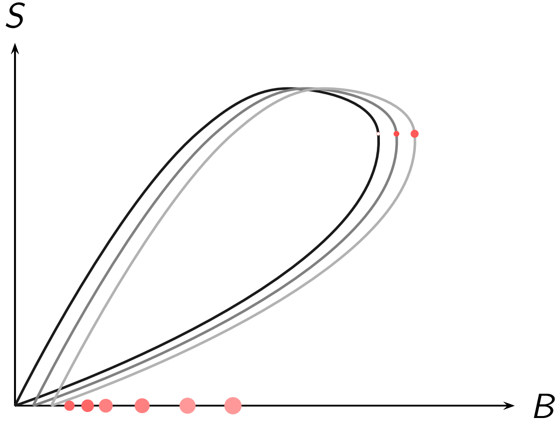

Buyer-Preferred Eq'm with Increasing Atom

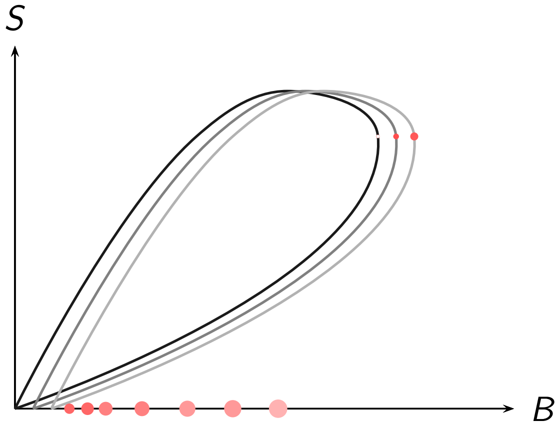

Buyer-Preferred Eq'm with Increasing Atom

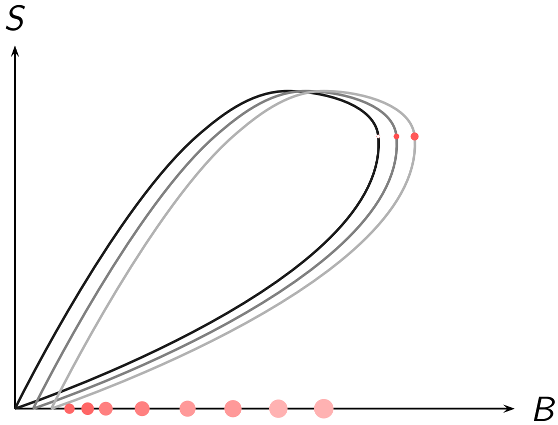

Buyer-Preferred Eq'm with Increasing Atom

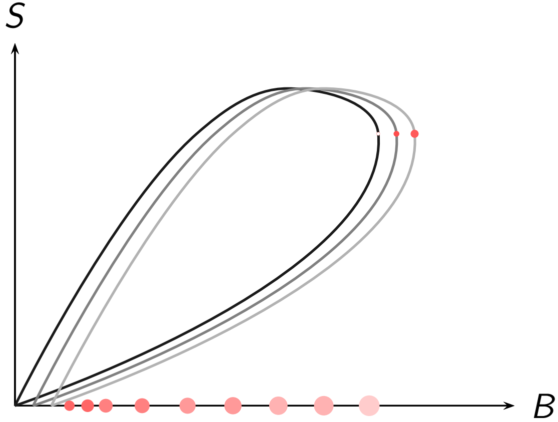

Buyer-Preferred Eq'm with Increasing Atom

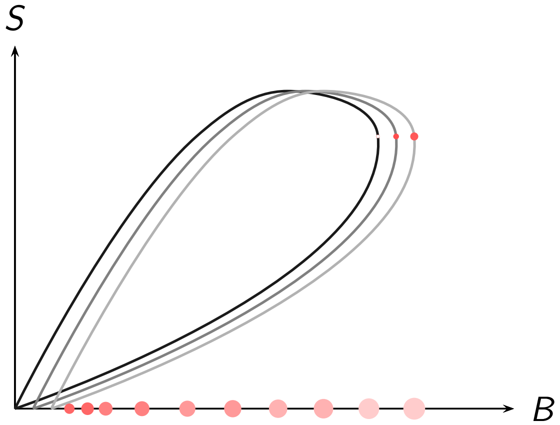

Buyer-Preferred Eq'm with Increasing Atom

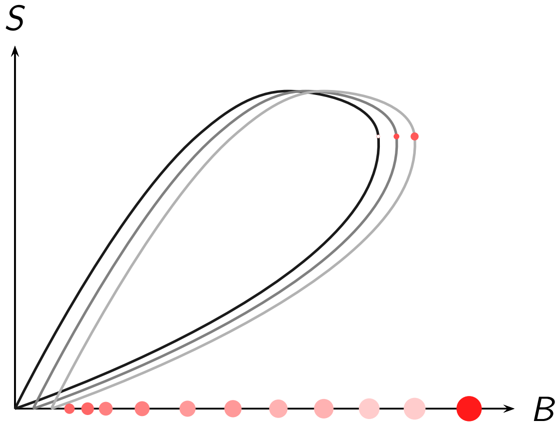

Related Literature

Repeated games with imperfect monitoring:

- Abreu, Pearce and Stacchetti (1990),
- Athey and Bagwell (2001).

"Trading favors:"

- Möbius (2001),
- Hauser and Hopenhayn (2008).

Relational contracts with random opportunity cost:

- Board (2011),
- Li and Matouschek (2013).

What Is Missing?

Thank You For Tuning In!

All is Observed: What if Prices Could Vary?

Buyer prefers to maximize overall surplus:

choose the efficient cutoff: $v^* = 1 - c$, and

give the seller enough rent to incentivize $q = 1$.

Given outside option v, buyer is willing to pay seller up to $1 - v$.

Seller is willing to set $q = 1$ if capturing the added surplus (*i.e.*, $(1 - v - c)$ whenever $v \leq 1 - c$) is enough.

All is Observed: What if the Buyer Picks One Price?

Buyer internalizes trade-off between the cost of

- coming "too often;"
- paying "too much."

Buyer chooses p such that $v^* = 1 - c$.