UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Summer 2020

Advanced Topics

Block 4
Jul 23 and 26, 2020

Game plan

[1] Food for thought
[2] Signaling
[3] Social learning
[4] Auctions
[5] Evolutionary game theory
[6] Repeated games (the prisoner's dilemma)
[7] Back to decision theory...
[8] MobLab

Food for thought

LUPI

Many players simultaneously chose an integer between 1 and 99,999. Whoever chooses the lowest unique positive integer (LUPI) wins.

Question What does an equilibrium model of behavior predict in this game?

The field version of LUPI, called Limbo, was introduced by the government-owned Swedish gambling monopoly Svenska Spel. Despite its complexity, there is a surprising degree of convergence toward equilibrium.

Games with population uncertainty relax the assumption that the exact number of players is common knowledge.

In particular, in a Poisson game (Myerson; 1998, 2000) the number of players N is a random variable that follows a Poisson distribution with mean n so the probability that $N = k$ is given by

$$\frac{e^{-n} n^k}{k!}$$

In the Swedish game the average number of players was $n = 53,783$ and number choices were positive integers up to $99,999$.

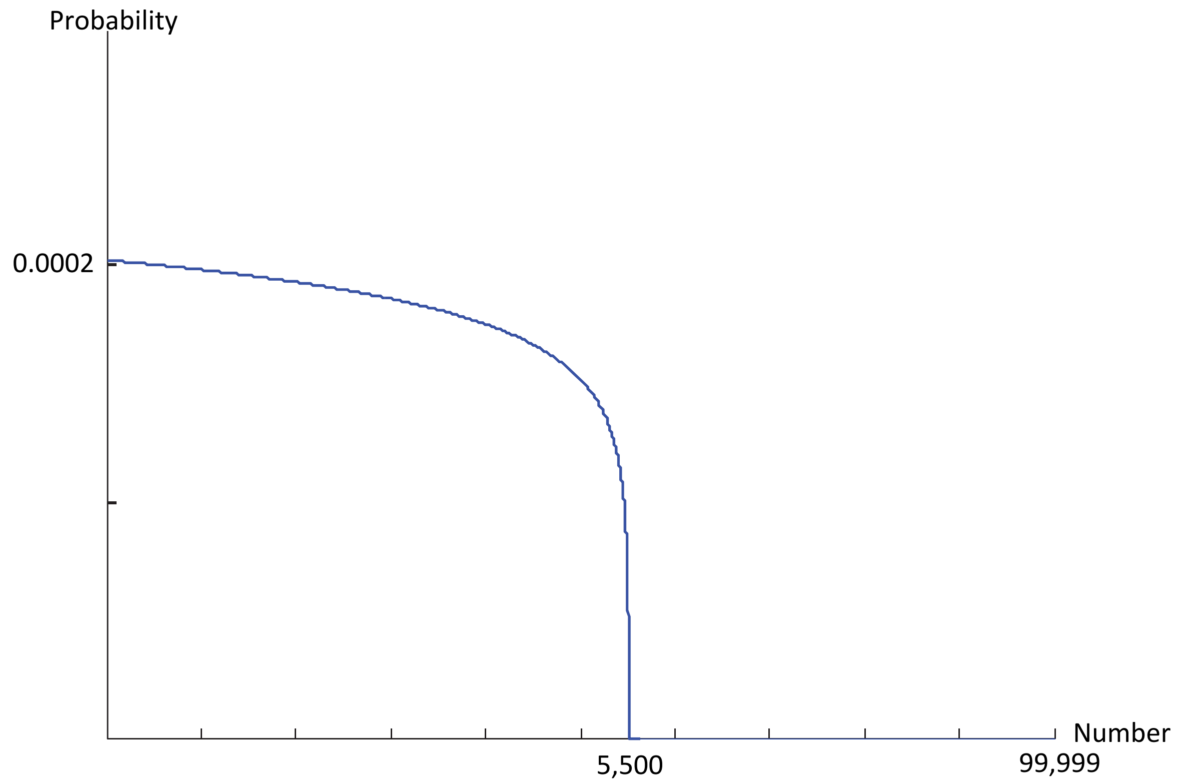
Probability
0.0002
Number
5,500
99,999

Morra

A two-player game in which each player simultaneously hold either one or two fingers and each guesses the total number of fingers held up.

If exactly one player guesses correctly, then the other player pays her the amount of her guess.

Question Model the situation as a strategic game and describe the equilibrium model of behavior predict in this game.

The game was played in ancient Rome, where it was known as “micatio.”

In Morra there are two players, each of whom has four (relevant) actions, S_1G_2, S_1G_3, S_2G_3, and S_2G_4, where S_iG_j denotes the strategy (Show i, Guess j).

The payoffs in the game are as follows

	S_1G_2	S_1G_3	S_2G_3	S_2G_4
S_1G_2	$0, 0$	$2, -2$	$-3, 3$	$0, 0$
S_1G_3	$-2, 2$	$0, 0$	$0, 0$	$3, -3$
S_2G_3	$3, -3$	$0, 0$	$0, 0$	$-4, 4$
S_2G_4	$0, 0$	$-3, 3$	$4, -4$	$0, 0$

Maximal game
(sealed-bid second-price auction)

Two bidders, each of whom privately observes a signal X_i that is independent and identically distributed (i.i.d.) from a uniform distribution on $[0, 10]$.

Let $X^{\mathsf{max}} = \mathsf{max}\{X_1, X_2\}$ and assume the ex-post common value to the bidders is X^{max}.

Bidders bid in a sealed-bid second-price auction where the highest bidder wins, earns the common value X^{max} and pays the second highest bid.

A review of the main ideas

We study two (out of four) groups of game theoretic models:

[1] Strategic games – all players simultaneously choose their plan of action once and for all.

[2] Extensive games (with perfect information) – players choose sequentially (and fully informed about all previous actions).

A solution (equilibrium) is a systematic description of the outcomes that may emerge in a family of games. We study two solution concepts:

[1] Nash equilibrium – a steady state of the play of a <u>strategic</u> game (no player has a profitable deviation given the actions of the other players).

[1] Subgame equilibrium – a steady state of the play of an <u>extensive</u> game (a Nash equilibrium in every <u>subgame</u> of the extensive game).

$\Longrightarrow$ Every subgame perfect equilibrium is also a Nash equilibrium.

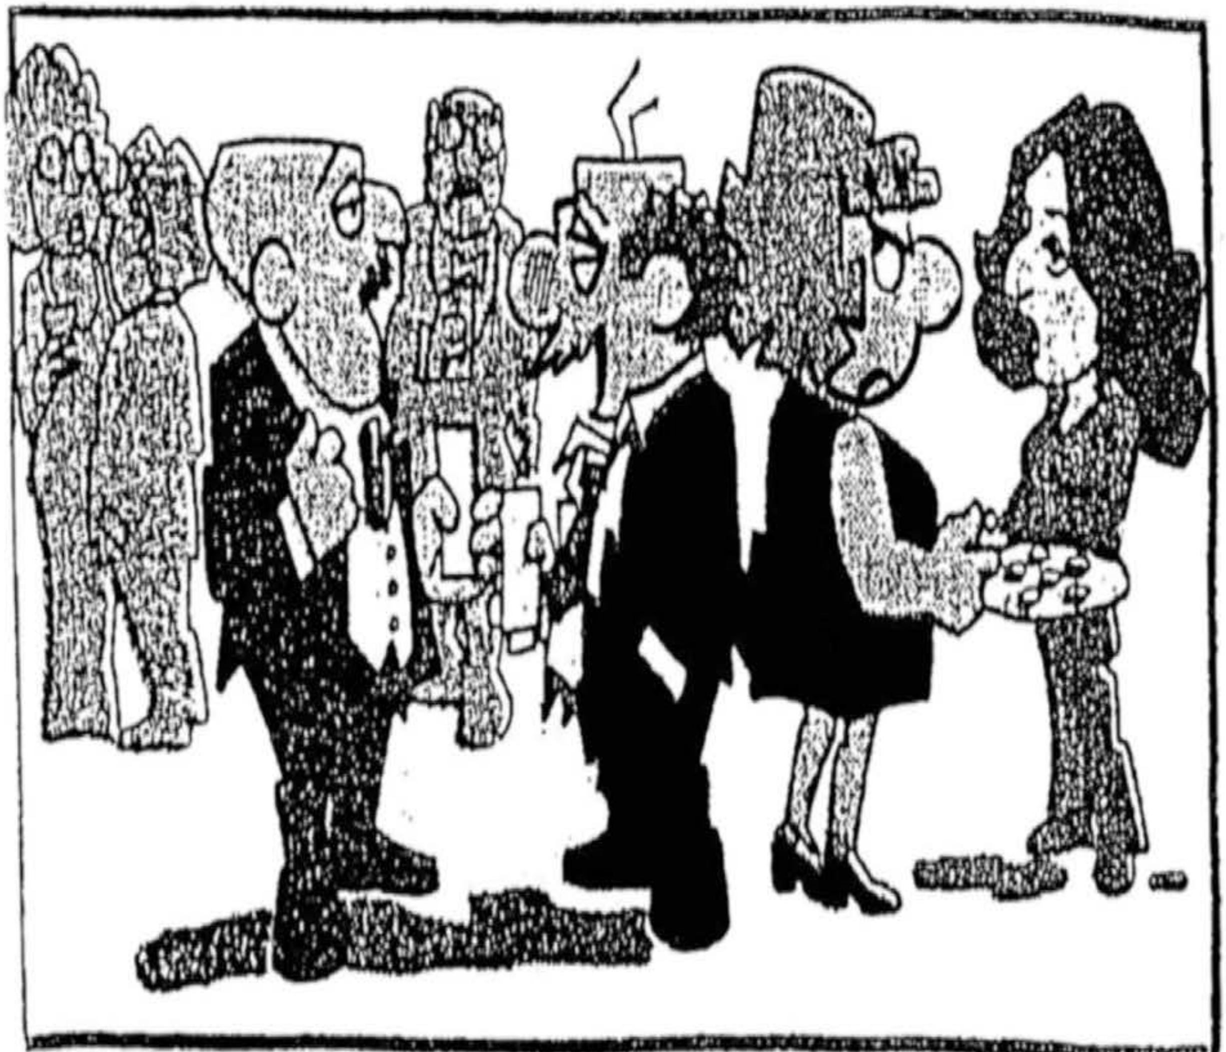

"LORETTA'S DRIVING BECAUSE I'M DRINKING, AND I'M DRINKING BECAUSE SHE'S DRIVING."

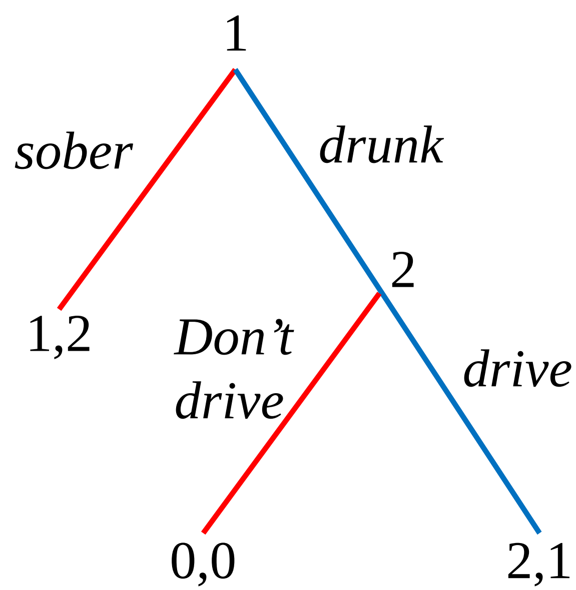
1
sober
drunk
2
1,2
Don't
drive
drive
0,0
2,1

Signaling
Spence's job-market signaling model

Signaling

- In the used-car market, owners of the good used cars have an incentive to try to convey the fact that they have a good car to the potential purchasers.
- Put differently, they would like choose actions that signal that they are offering a plum rather than a lemon.
- In some case, the presence of a "signal" allows the market to function more effectively than it would otherwise.

Example – educational signaling

– Suppose that a fraction $0 < b < 1$ of workers are *competent* and a fraction $1 - b$ are *incompetent*.

– The competent workers have marginal product of a_2 and the incompetent have marginal product of $a_1 < a_2$.

– For simplicity we assume a competitive labor market and a linear production function

$$L_1 a_1 + L_2 a_2$$

where L_1 and L_2 is the number of incompetent and competent workers, respectively.

- If worker quality is observable, then firm would just offer wages

$$w_1 = a_1 \text{ and } w_2 = a_2$$

  to competent workers, respectively.

- That is, each worker will paid his marginal product and we would have an <u>efficient</u> equilibrium.

- But what if the firm cannot observe the marginal products so it cannot distinguish the two types of workers?

- If worker quality is unobservable, then the "best" the firm can do is to offer the average wage

$$w = (1 - b)a_1 + ba_2.$$

- If both types of workers agree to work at this wage, then there is no problem with adverse selection (more below).
- The incompetent (resp. competent) workers are getting paid more (resp. less) than their marginal product.

- The competent workers would like a way to signal that they are more productive than the others.

- Suppose now that there is some signal that the workers can acquire that will distinguish the two types

- One nice example is education – it is cheaper for the competent workers to acquire education than the incompetent workers.

- To be explicit, suppose that the cost (dollar costs, opportunity costs, costs of the effort, etc.) to acquiring e years of education is

$$c_1 e \text{ and } c_2 e$$

  for incompetent and competent workers, respectively, where $c_1 > c_2$.

- Suppose that workers conjecture that firms will pay a wage $s(e)$ where s is some increasing function of e.

- Although education has no effect on productivity (MBA?), firms may still find it profitable to base wage on education – attract a higher-quality work force.

Market equilibrium

In the educational signaling example, there appear to be several possibilities for equilibrium:

[1] The (representative) firm offers a single contract that attracts both types of workers.

[2] The (representative) firm offers a single contract that attracts only one type of workers.

[3] The (representative) firm offers two contracts, one for each type of workers.

- A separating equilibrium involves each type of worker making a choice that separate himself from the other type.

- In a pooling equilibrium, in contrast, each type of workers makes the same choice, and all getting paid the wage based on their average ability.

  Note that a separating equilibrium is wasteful in a social sense – no social gains from education since it does not change productivity.

Example (cont.)

– Let e_1 and e_2 be the education level actually chosen by the workers. Then, a separating (signaling) equilibrium has to satisfy:

[1] zero-profit conditions

$$\begin{aligned} s(e_1) &= a_1 \\ s(e_2) &= a_2 \end{aligned}$$

[2] self-selection conditions

$$\begin{aligned} s(e_1) - c_1 e_1 &\geq s(e_2) - c_1 e_2 \\ s(e_2) - c_2 e_2 &\geq s(e_1) - c_2 e_1 \end{aligned}$$

– In general, there may by many functions $s(e)$ that satisfy conditions [1] and [2]. One wage profile consistent with separating equilibrium is

$$s(e) = \begin{cases} a_2 & \text{if} \quad e > e^* \\ a_1 & \text{if} \quad e \leq e^* \end{cases}$$

and

$$\frac{a_2 - a_1}{c_2} > e^* > \frac{a_2 - a_1}{c_1}$$

$\Longrightarrow$ Signaling can make things better or worse – each case has to examined on its own merits!

The Sheepskin (diploma) effect

The increase in wages associated with obtaining a higher credential:

- Graduating high school increases earnings by 5 to 6 times as much as does completing a year in high school that does not result in graduation.
- The same discontinuous jump occurs for people who graduate from collage.
- High school graduates produce essentially the same amount of output as non-graduates.

Social Learning
Herd behavior and informational cascades

"Men nearly always follow the tracks made by others and proceed in their affairs by imitation." Machiavelli (Renaissance philosopher)

Examples

Business strategy

- TV networks make introductions in the same categories as their rivals.

Finance

- The withdrawal behavior of small number of depositors starts a bank run.

Politics

– The solid New Hampshirites (probably) can not be too far wrong.

Crime

– In NYC, individuals are more likely to commit crimes when those around them do.

Why should individuals behave in this way?

Several "theories" explain the existence of uniform social behavior:

– benefits from conformity

– sanctions imposed on deviants

– network / payoff externalities

– social learning

Broad definition: any situation in which individuals learn by observing the behavior of others.

Informational cascades and herd behavior

Two phenomena that have elicited particular interest are *informational cascades* and *herd behavior*.

– Cascade: agents 'ignore' their private information when choosing an action.

– Herd: agents choose the same action, not necessarily ignoring their private information.

- While the terms informational cascade and herd behavior are used interchangeably there is a significant difference between them.

- In an informational cascade, an agent considers it optimal to follow the behavior of her predecessors without regard to her private signal.

- When acting in a herd, agents choose the same action, not necessarily ignoring their private information.

- Thus, an informational cascade implies a herd but a herd is not necessarily the result of an informational cascade.

A model of social learning

Signals

- Each player $n \in \{1, ..., N\}$ receives a signal θ_n that is private information.

- For simplicity, $\{\theta_n\}$ are independent and uniformly distributed on $[-1, 1]$.

Actions

- Sequentially, each player n has to make a binary irreversible decision $x_n \in \{0, 1\}$.

Payoffs

– $x = 1$ is profitable if and only if $\sum_{n \leq N} \theta_n \geq 0$, and $x = 0$ is profitable otherwise.

Information

– Perfect information

$$\mathcal{I}_n = \{\theta_n, (x_1, x_2, ..., x_{n-1})\}$$

– Imperfect information

$$\mathcal{I}_n = \{\theta_n, x_{n-1}\}$$

Sequential social-learning model:
Well heck, if all you smart cookies agree, who am I to dissent?

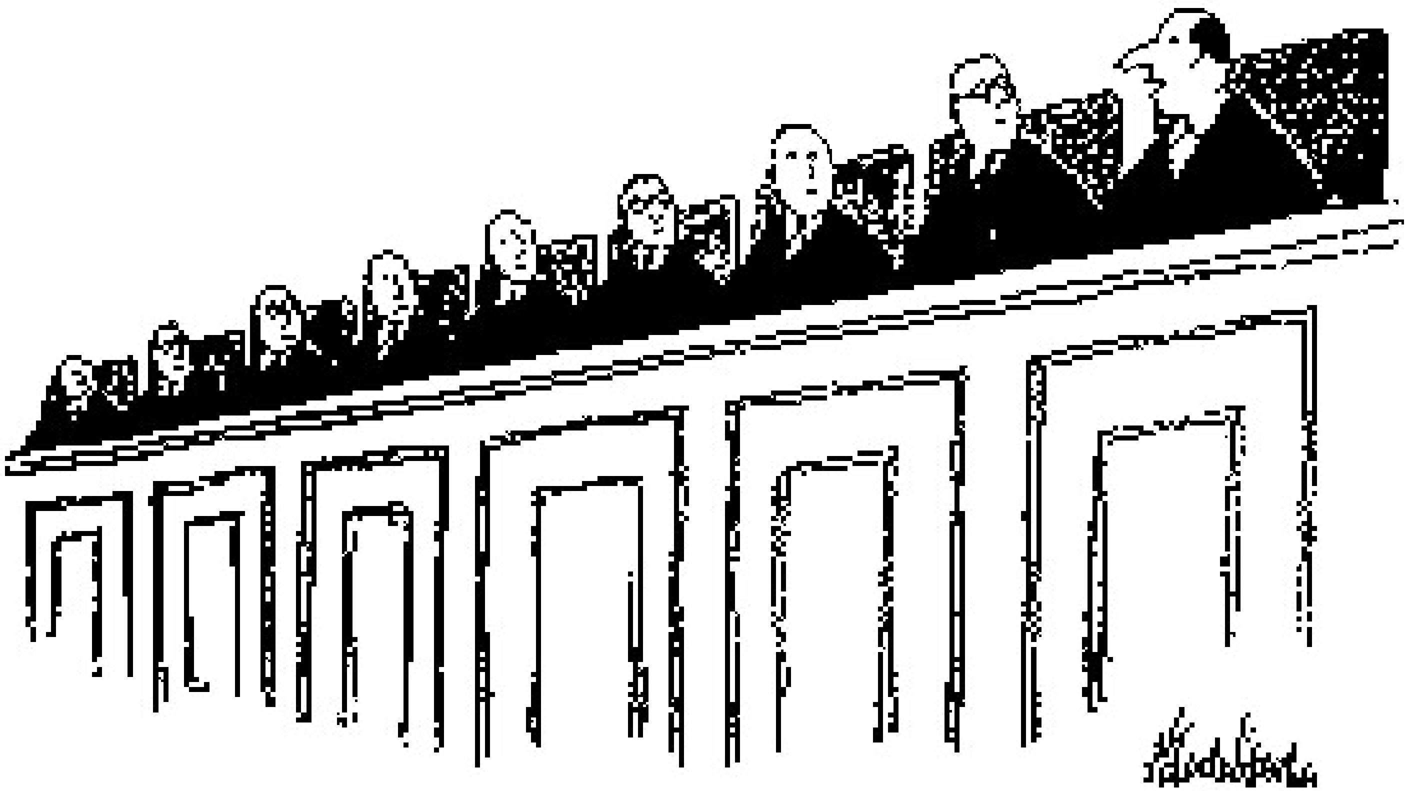

Imperfect information:
Which way is the wind blowing?!

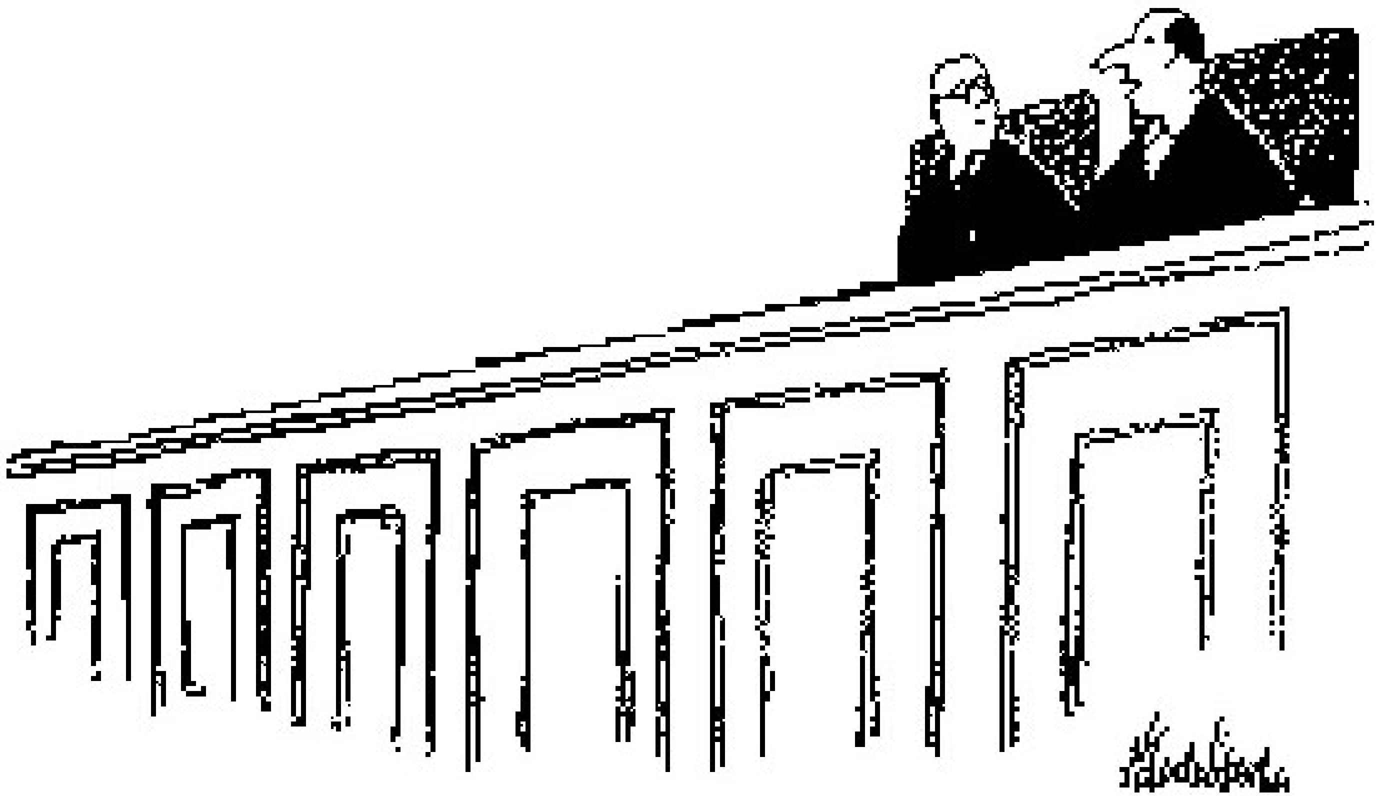

A three-agent example

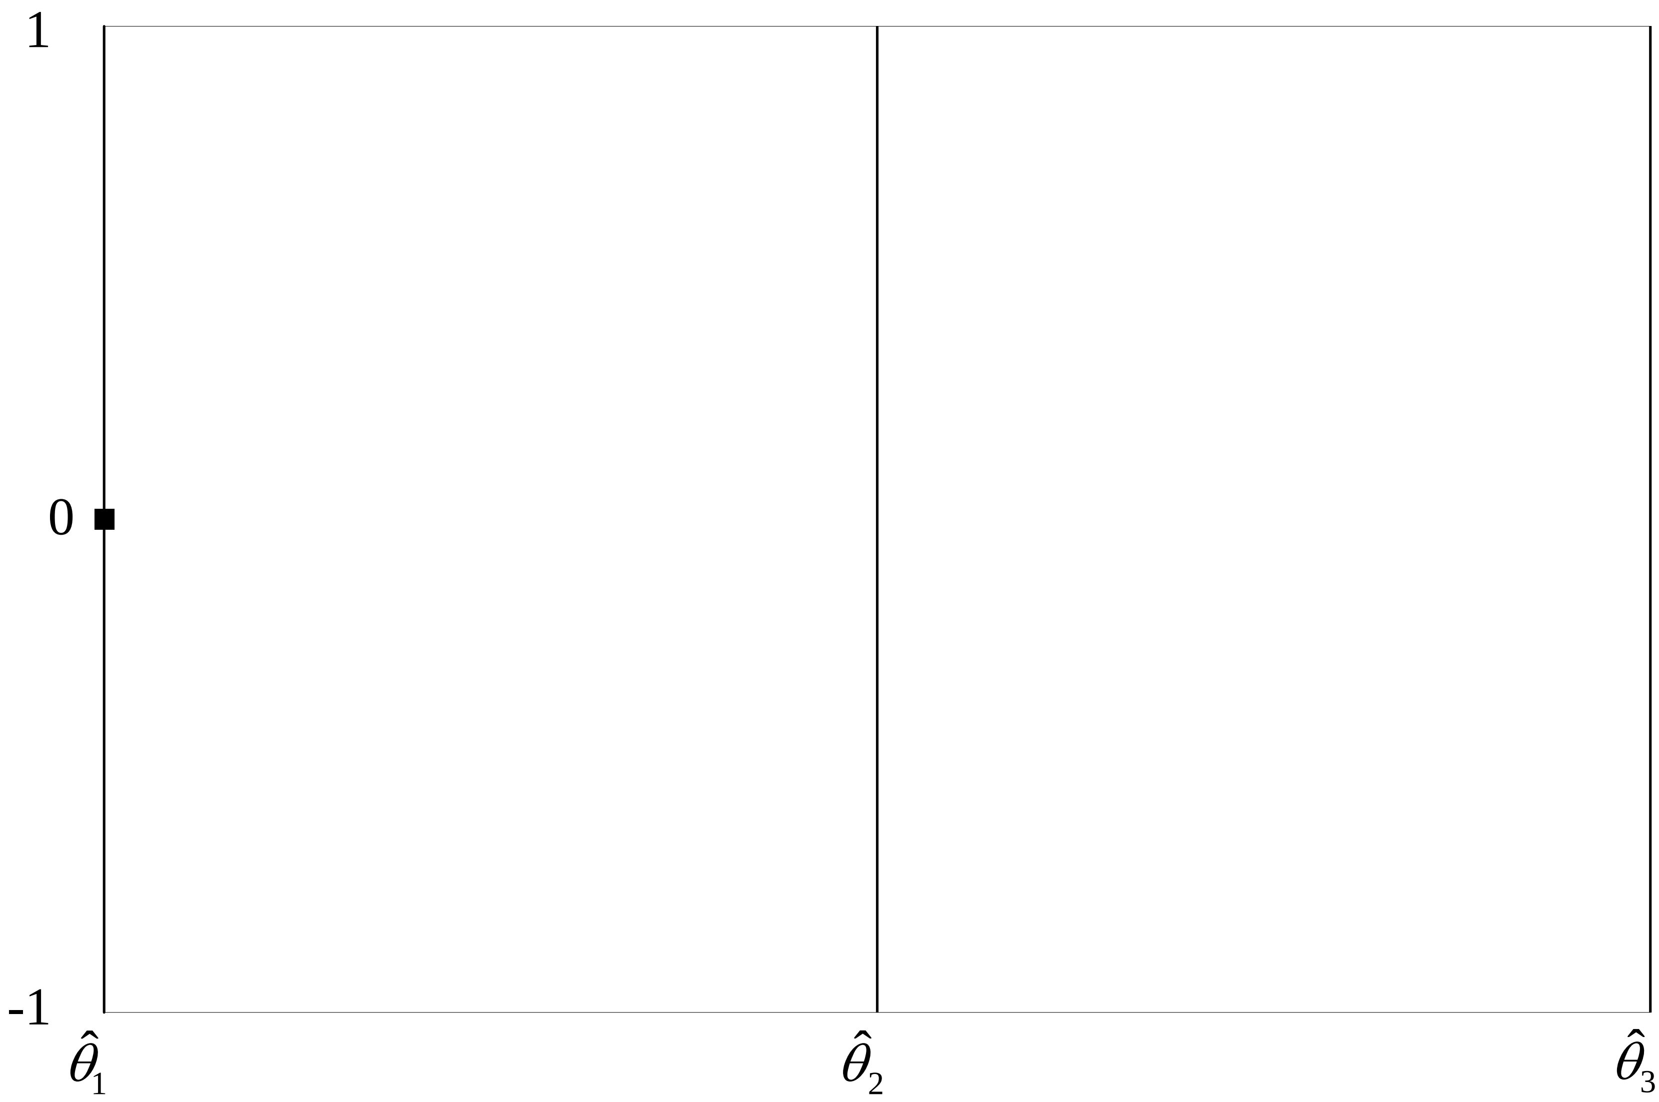

A three-agent example

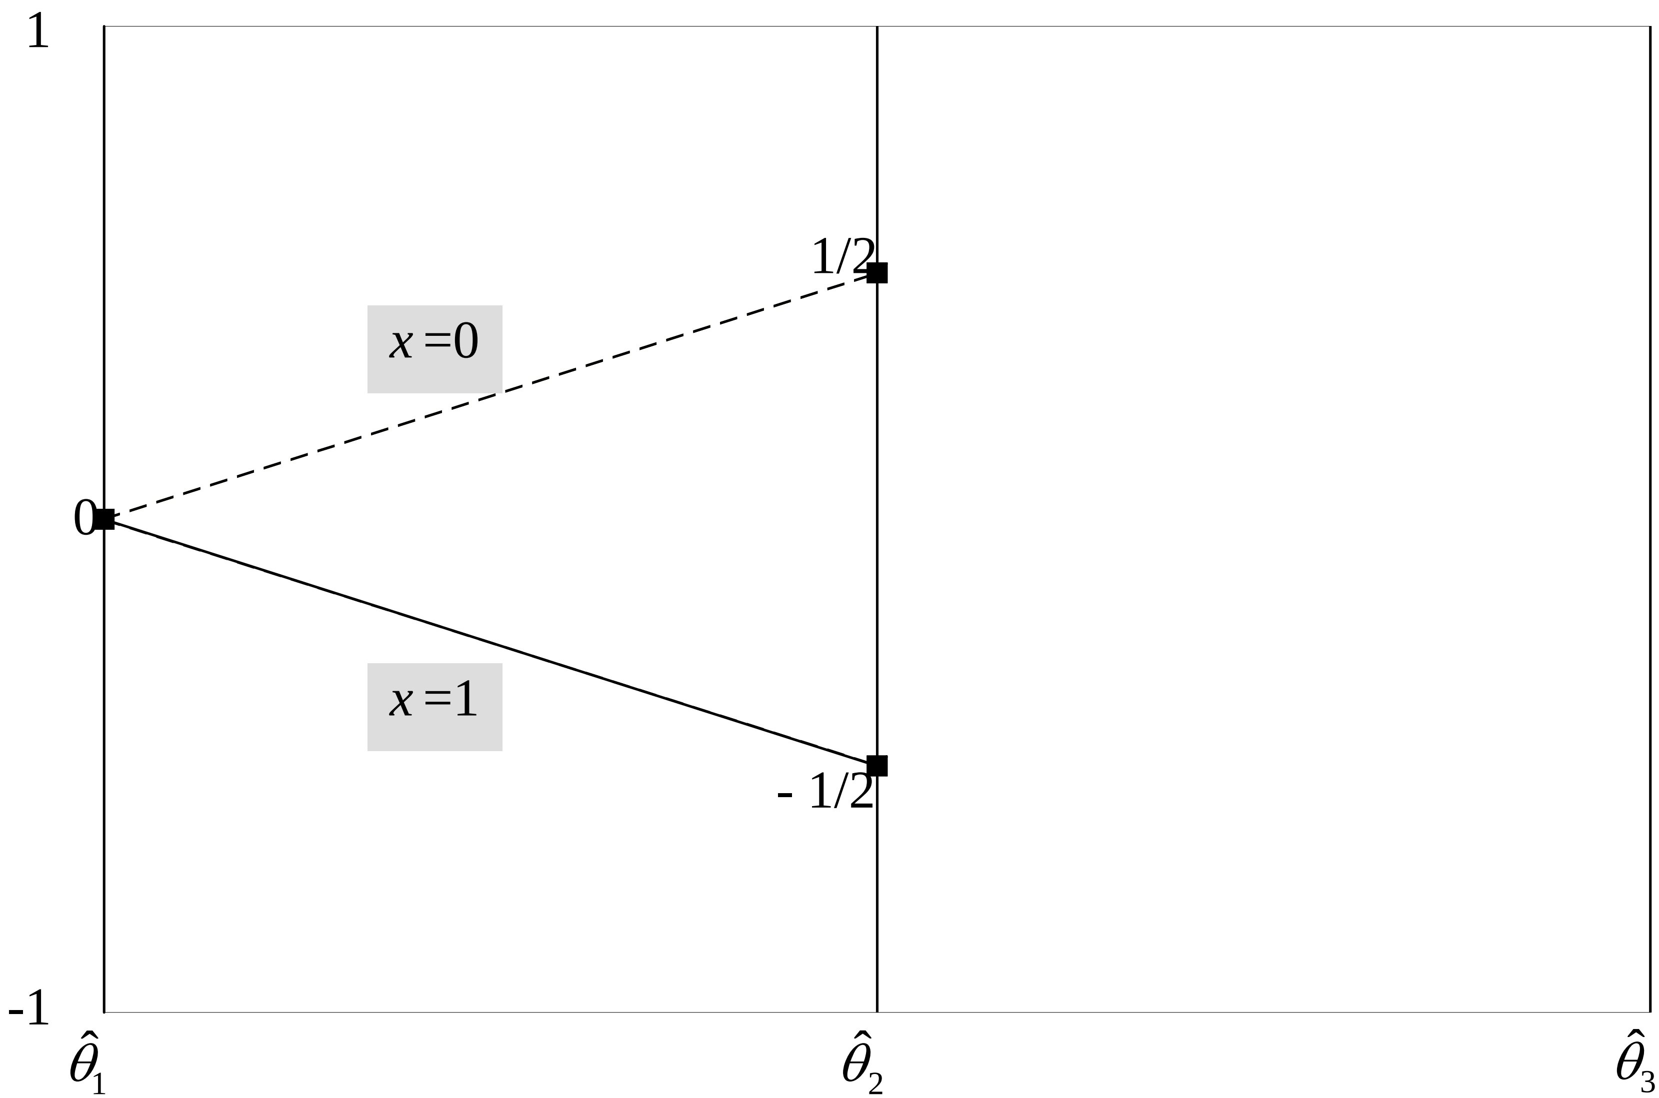

A three-agent example under perfect information

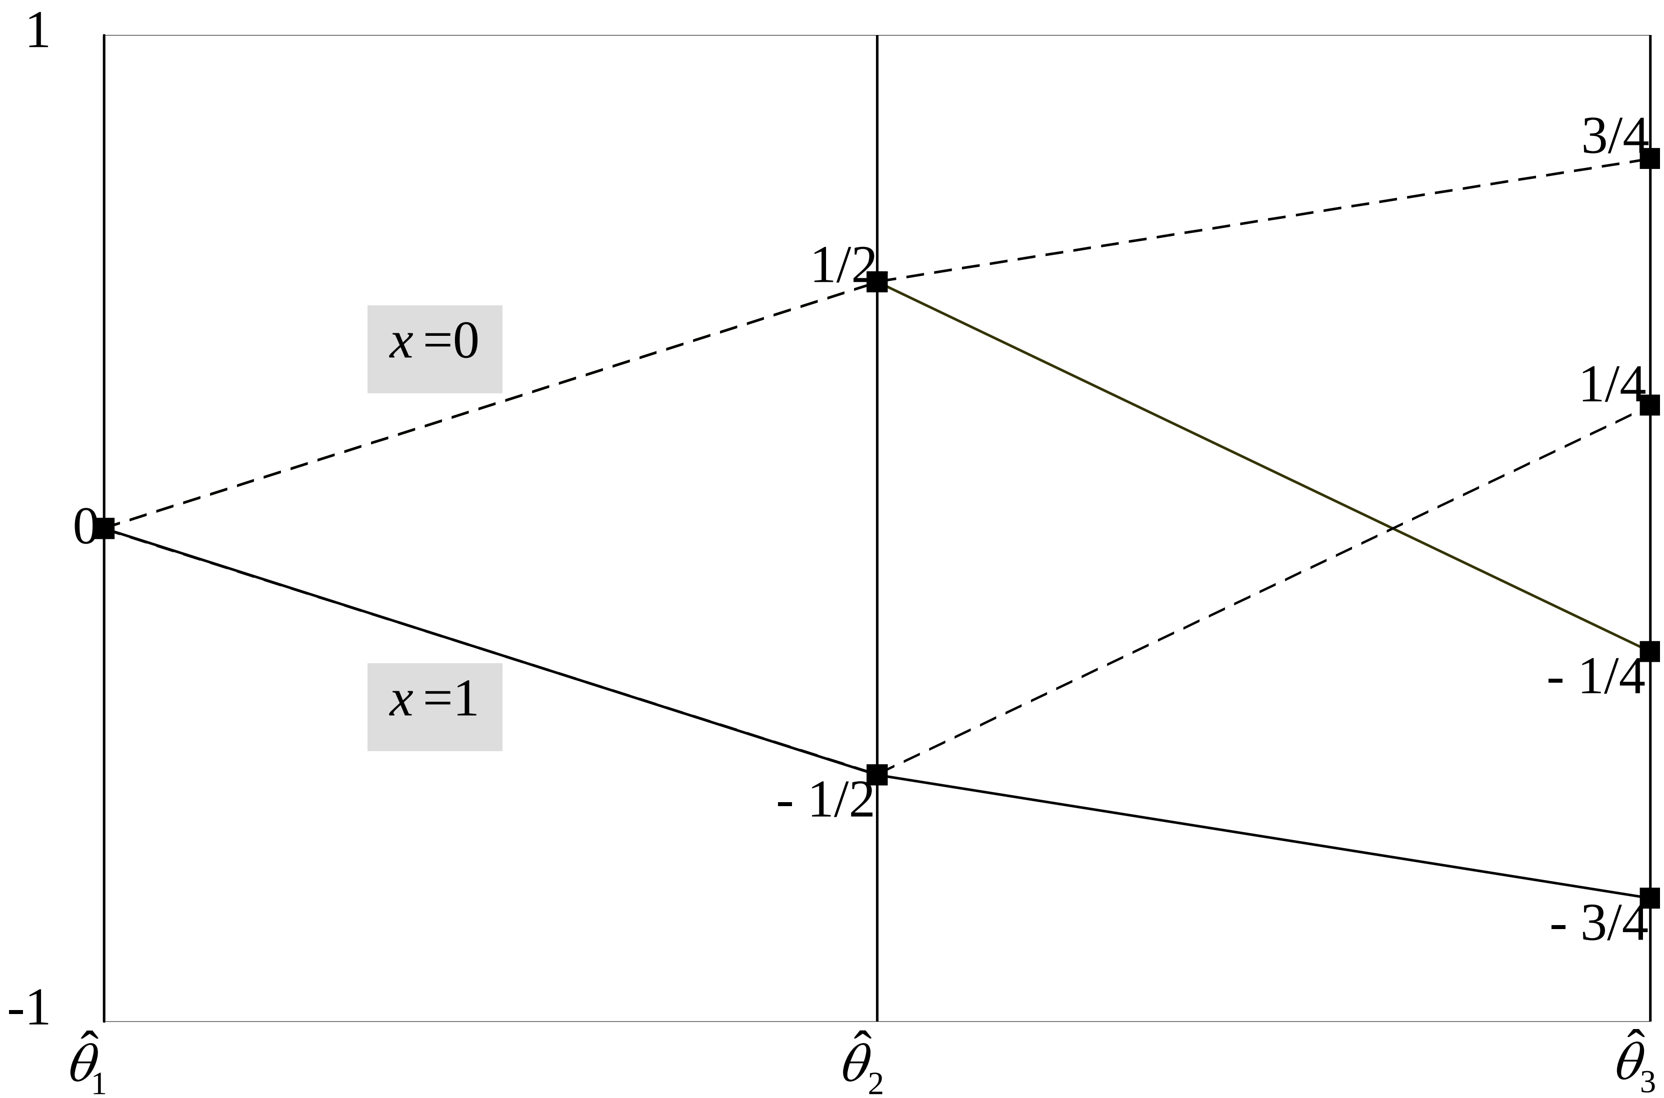

A three-agent example under imperfect information

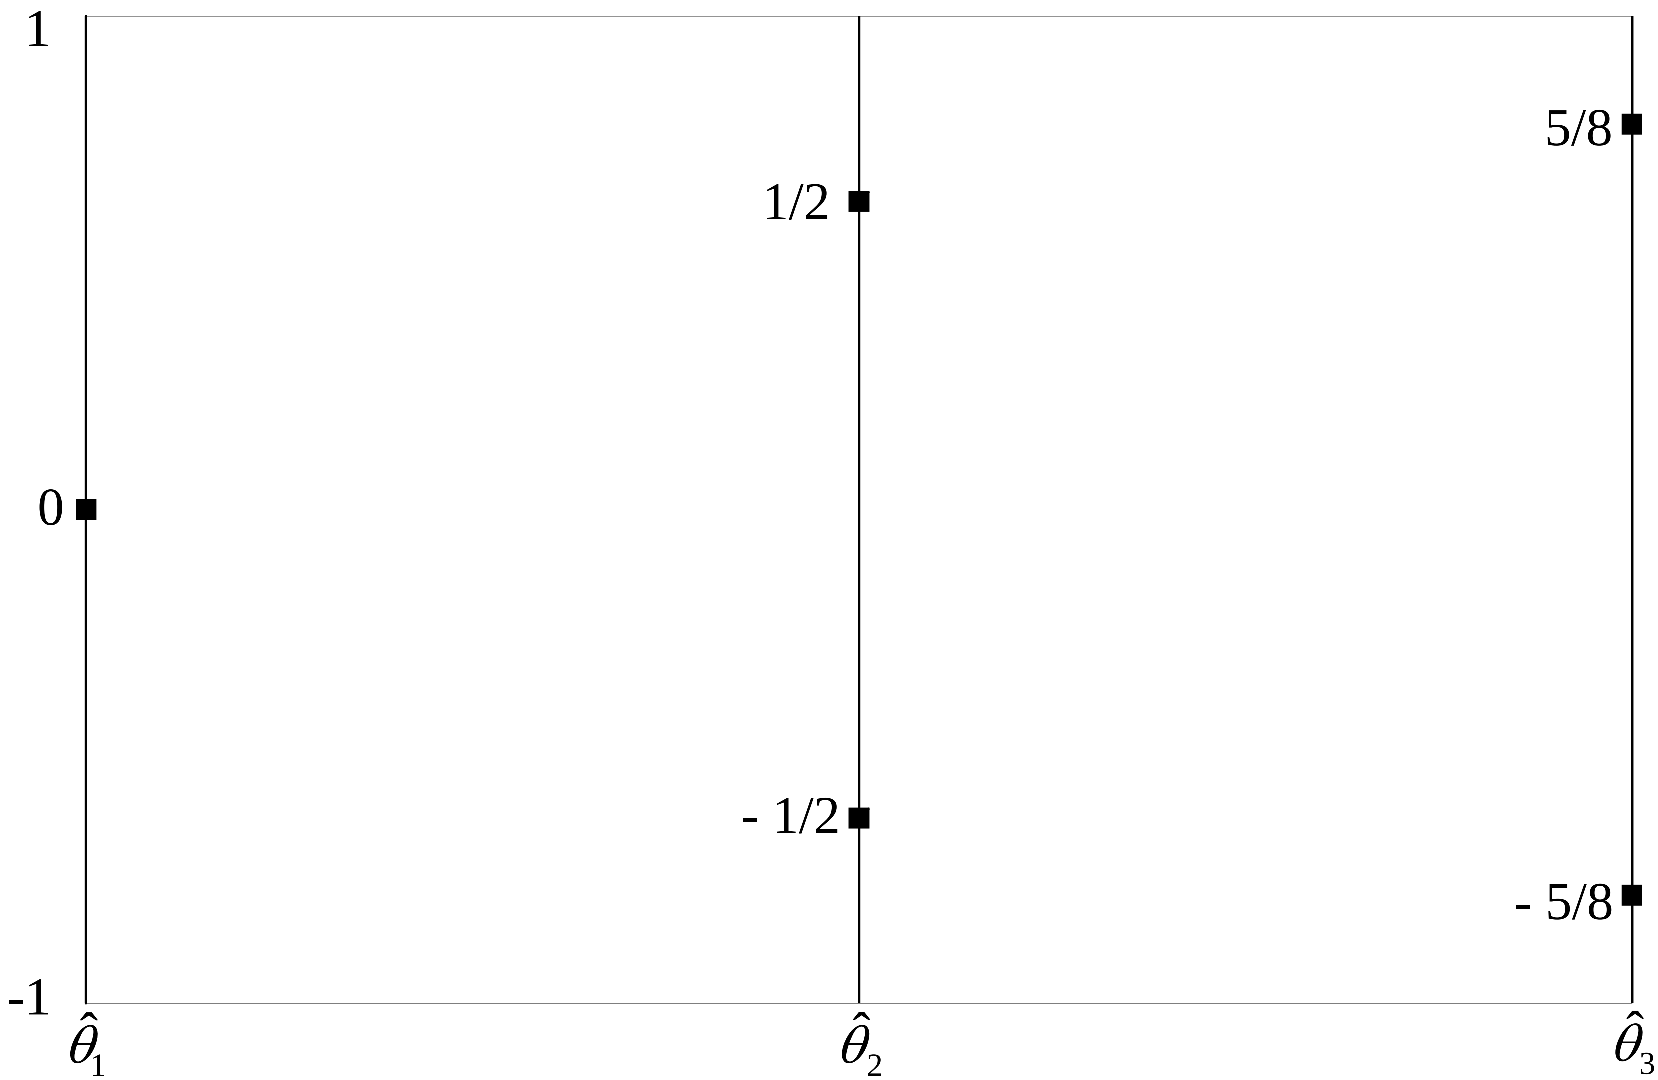

A sequence of cutoffs under imperfect and perfect information

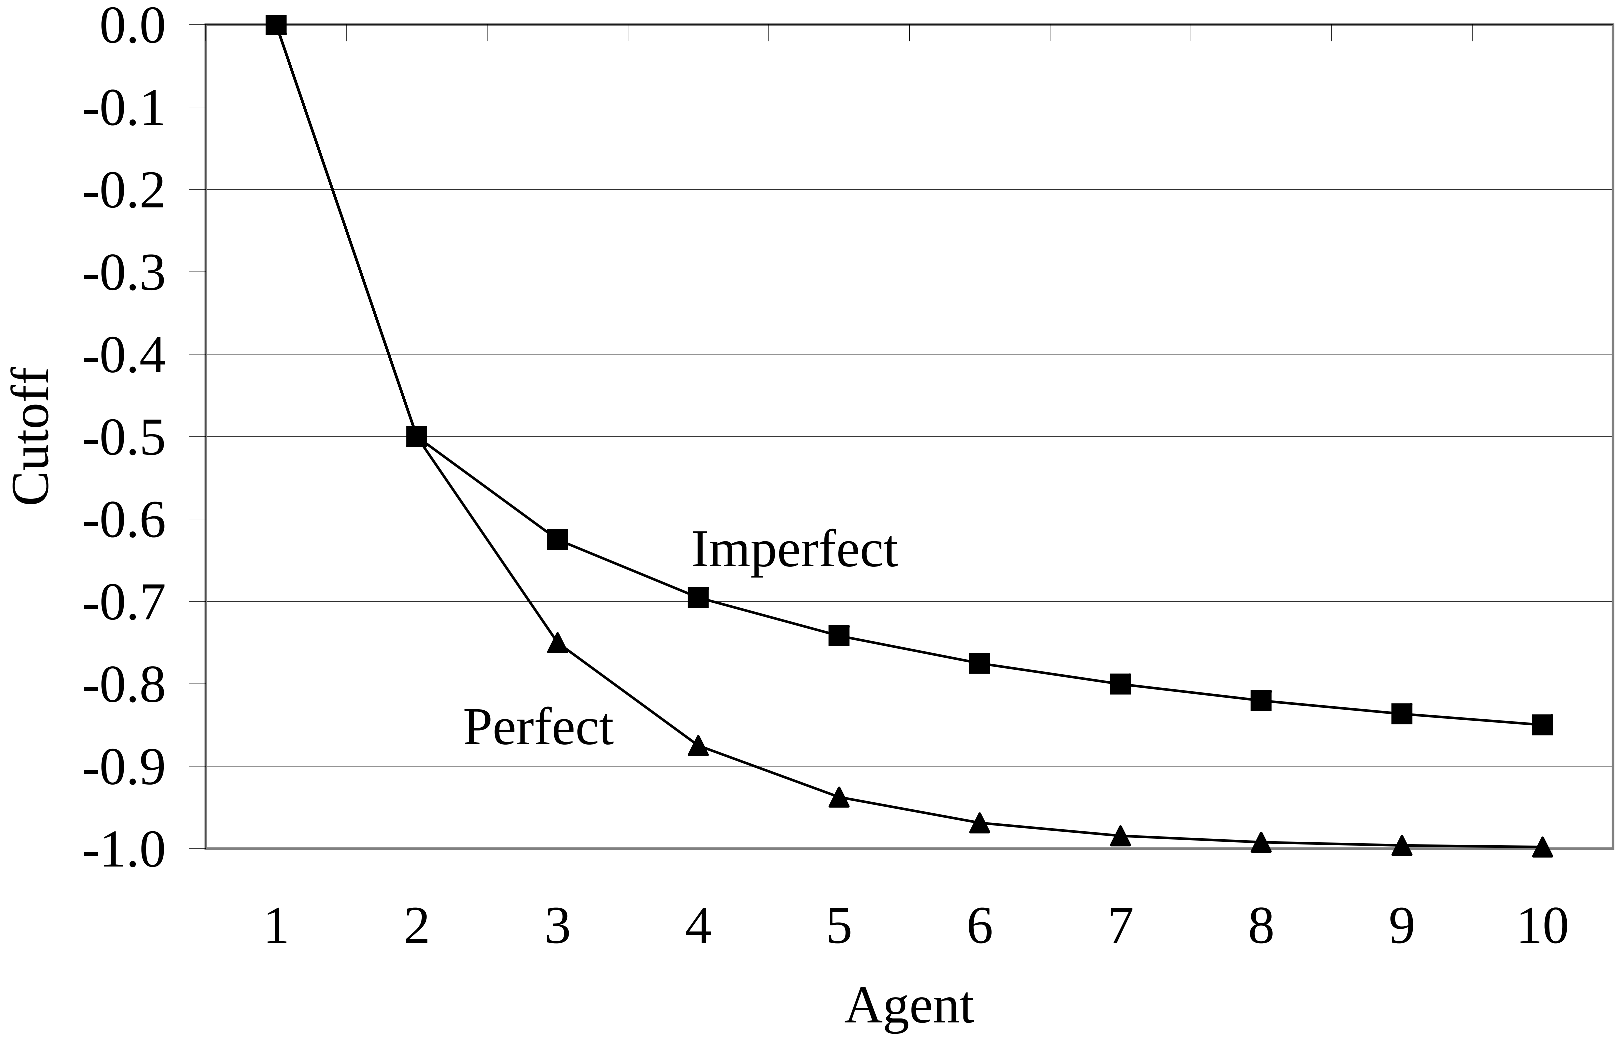

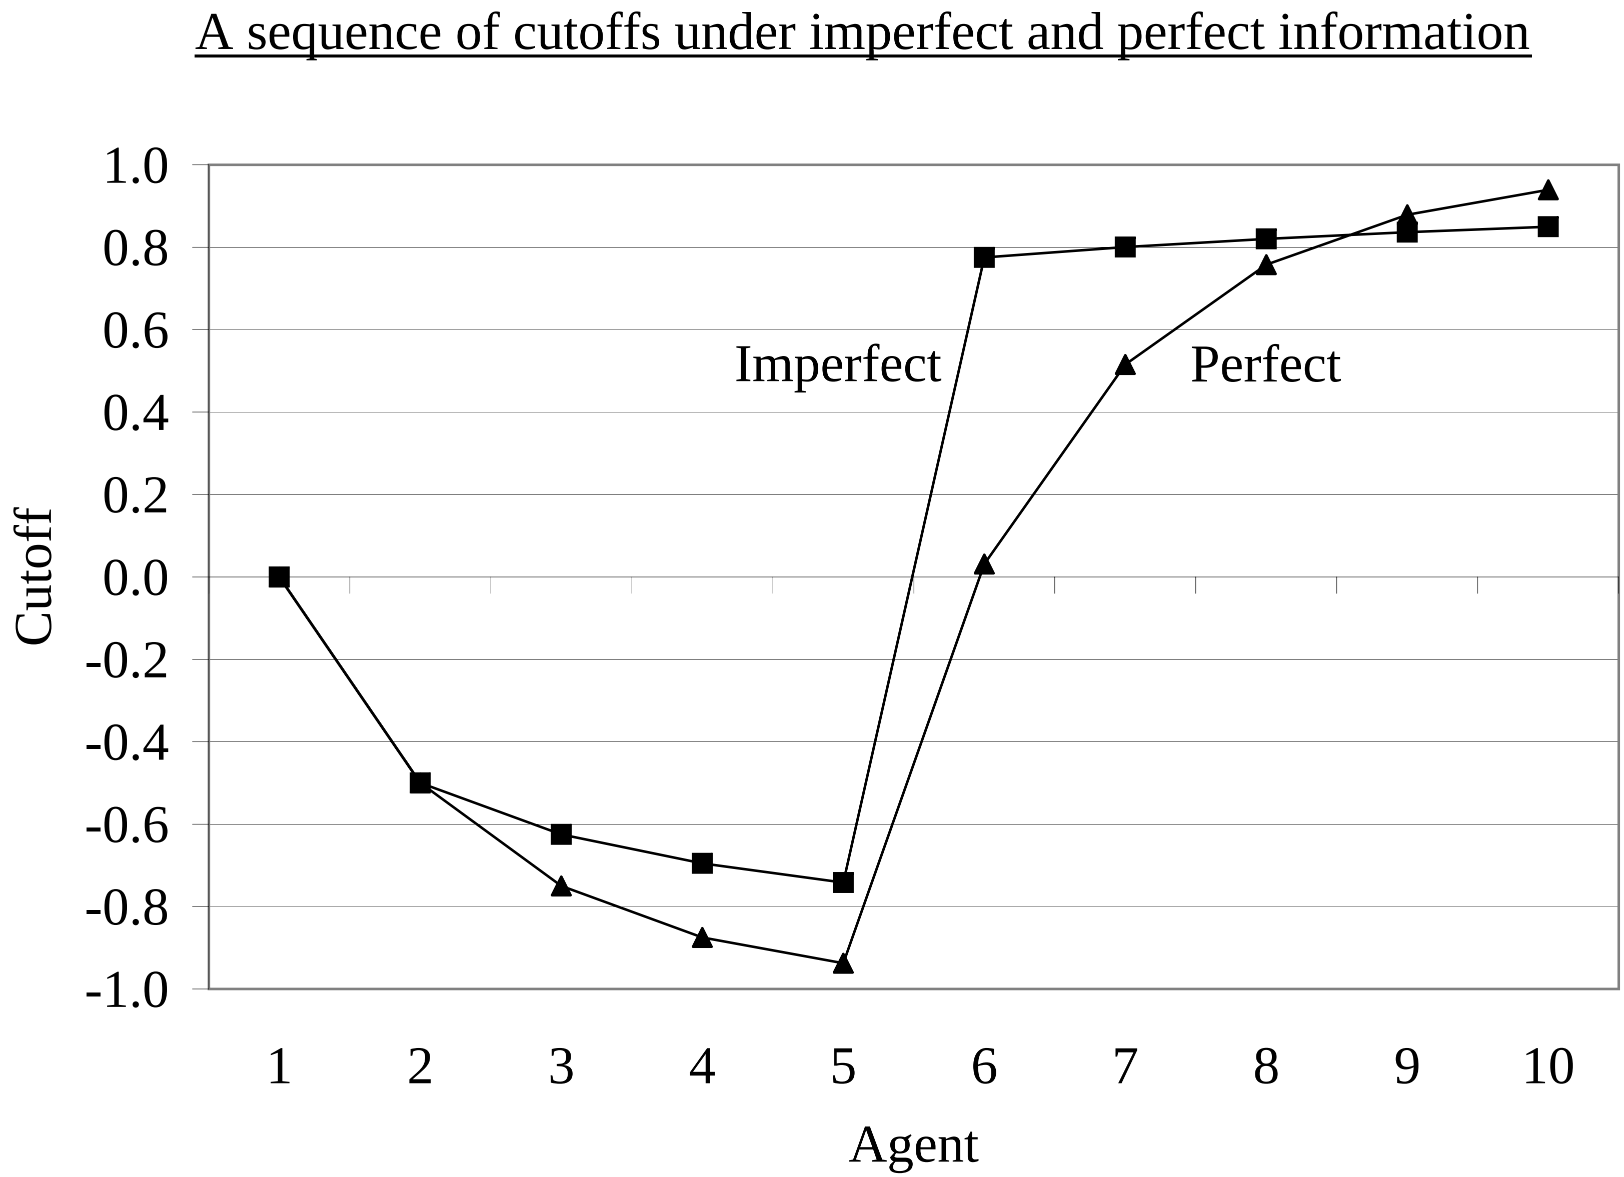
A sequence of cutoffs under imperfect and perfect information
Cutoff
1.0
0.8
0.6
0.4
0.2
0.0
-0.2
-0.4
-0.6
-0.8
-1.0
Imperfect
Perfect
1
2
3
4
5
6
7
8
9
10
Agent

The decision problem

– The optimal decision rule is given by

$$x_n = 1 \text{ if and only if } \mathbb{E}\left[\sum_{i=1}^{N} \theta_i \mid \mathcal{I}_n\right] \geq 0.$$

Since $\mathcal{I}_n$ does not provide any information about the content of successors' signals, we obtain

$$x_n = 1 \text{ if and only if } \mathbb{E}\left[\sum_{i=1}^{n} \theta_i \mid \mathcal{I}_n\right] \geq 0$$

Hence,

$$x_n = 1 \text{ if and only if } \theta_n \geq -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid \mathcal{I}_n\right].$$

The cutoff process

– For any n, the optimal strategy is the *cutoff strategy*

$$x_n = \begin{cases} 1 & if \quad \theta_n \geq \hat{\theta}_n \\ 0 & if \quad \theta_n < \hat{\theta}_n \end{cases}$$

where

$$\hat{\theta}_n = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid \mathcal{I}_n\right]$$

is the optimal history-contingent cutoff.

– $\hat{\theta}_n$ is sufficient to characterize the individual behavior, and $\{\hat{\theta}_n\}$ characterizes the social behavior of the economy.

Overview of results

Perfect information

- A cascade need not arise, but herd behavior must arise.

Imperfect information

- Herd behavior is impossible. There are periods of uniform behavior, punctuated by increasingly rare switches.

- The similarity:
  - Agents can, for a long time, make the same (incorrect) choice.
- The difference:
  - Under perfect information, a herd is an absorbing state. Under imperfect information, continued, occasional and sharp shifts in behavior.

- The dynamics of social learning depend crucially on the extensive form of the game.

- The key economic phenomenon that imperfect information captures is a succession of fads starting suddenly, expiring rather easily, each replaced by another fad.

- The kind of episodic instability that is characteristic of socioeconomic behavior in the real world makes more sense in the imperfect-information model.

As such, the imperfect-information model gives insight into phenomena such as manias, fashions, crashes and booms, and better answers such questions as:

– Why do markets move from boom to crash without settling down?

– Why is a technology adopted by a wide range of users more rapidly than expected and then, suddenly, replaced by an alternative?

– What makes a restaurant fashionable over night and equally unexpectedly unfashionable, while another becomes the ‘in place’, and so on?

The case of perfect information

The optimal history-contingent cutoff rule is

$$\hat{\theta}_n = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_1, ..., x_{n-1}\right],$$

and $\hat{\theta}_n$ is different from $\hat{\theta}_{n-1}$ only by the information reveals by the action of agent $(n-1)$

$$\hat{\theta}_n = \hat{\theta}_{n-1} - \mathbb{E}\left[\theta_{n-1} \mid \hat{\theta}_{n-1}, x_{n-1}\right],$$

The cutoff dynamics thus follow the cutoff process

$$\hat{\theta}_n = \begin{cases} \frac{-1+\hat{\theta}_{n-1}}{2} & \text{if} \quad x_{n-1} = 1 \\ \frac{1+\hat{\theta}_{n-1}}{2} & \text{if} \quad x_{n-1} = 0 \end{cases}$$

where $\hat{\theta}_1 = 0$.

Informational cascades

– $-1 < \hat{\theta}_n < 1$ for any n so any player takes his private signal into account in a non-trivial way.

Herd behavior

– $\{\hat{\theta}_n\}$ has the martingale property by the Martingale Convergence Theorem a limit-cascade implies a herd.

The case of imperfect information

The optimal history-contingent cutoff rule is

$$\hat{\theta}_n = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1}\right],$$

which can take two values conditional on $x_{n-1} = 1$ or $x_{n-1} = 0$

$$\begin{aligned} \overline{\theta}_n &= -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1} = 1\right], \\ \underline{\theta}_n &= -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1} = 1\right]. \end{aligned}$$

where $\overline{\theta}_n = -\underline{\theta}_n$.

The law of motion for $\overline{\theta}_n$ is given by

$$\begin{aligned}\overline{\theta}_n = P(x_{n-2} = 1|x_{n-1} = 1)\left\{\overline{\theta}_{n-1} - \mathbb{E}\left[\theta_{n-1} \mid x_{n-2} = 1\right]\right\} \\ + P(x_{n-2} = 0|x_{n-1} = 1)\left\{\underline{\theta}_{n-1} - \mathbb{E}\left[\theta_{n-1} \mid x_{n-2} = 0\right]\right\},\end{aligned}$$

which simplifies to

$$\begin{aligned}\overline{\theta}_n &= \frac{1 - \overline{\theta}_{n-1}}{2}\left[\overline{\theta}_{n-1} - \frac{1 + \overline{\theta}_{n-1}}{2}\right] \\ &\quad + \frac{1 - \underline{\theta}_{n-1}}{2}\left[\underline{\theta}_{n-1} - \frac{1 + \underline{\theta}_{n-1}}{2}\right].\end{aligned}$$

Given that $\overline{\theta}_n = -\overline{\theta}_n$, the cutoff dynamics under imperfect information follow the cutoff process

$$\hat{\theta}_n = \begin{cases} -\frac{1+\hat{\theta}_{n-1}^2}{2} & \text{if} \quad x_{n-1} = 1 \\ \frac{1+\hat{\theta}_{n-1}^2}{2} & \text{if} \quad x_{n-1} = 0 \end{cases}$$

where $\hat{\theta}_1 = 0$.

Informational cascades

– $-1 < \hat{\theta}_n < 1$ for any n so any player takes his private signal into account in a non-trivial way.

Herd behavior

– $\{\hat{\theta}_n\}$ is not convergent (proof is hard!) and the divergence of cutoffs implies divergence of actions.

– Behavior exhibits periods of uniform behavior, punctuated by increasingly rare switches.

Auctions

Auction design

Two important issues for auction design are:

- Attracting entry
- Preventing collusion

Sealed-bid auction deals better with these issues, but it is more likely to lead to inefficient outcomes.

European 3G mobile telecommunication license auctions

Although the blocks of spectrum sold were very similar across countries, there was an enormous variation in revenues (in USD) per capita:

Austria	100
Belgium	45
Denmark	95
Germany	615
Greece	45
Italy	240
Netherlands	170
Switzerland	20
United Kingdom	650

United Kingdom

- 4 licenses to be auctioned off and 4 incumbents (with advantages in terms of costs and brand).

- To attract entry and deter collusion – an English until 5 bidders remain and then a sealed-bid with reserve price given by lowest bid in the English.

- later a 5th license became available to auction, a straightforward English auction was implemented.

Netherlands

– Followed UK and used (only) an English auction, but they had 5 incumbents and 5 licenses!

– Low participation: strongest potential entrants made deals with incumbents, and weak entrants either stayed out or quit bidding.

Switzerland

- Also followed UK and ran an English auction for 4 licenses. Companies either stayed out or quit bidding.

  1. The government permitted last-minute joint-bidding agreements. Demand shrank from 9 to 4 bidders one week before the auction.

  2. The reserve price had been set too low. The government tried to change the rules but was opposed by remaining bidders and legally obliged to stick to the original rules.

- Collected 1/30 per capita of UK, and 1/50 of what they had hoped for!

Types of auctions

Sequential / simultaneous

Bids may be called out sequentially or may be submitted simultaneously in sealed envelopes:

- English (or oral) – the seller actively solicits progressively higher bids and the item is soled to the highest bidder.
- Dutch – the seller begins by offering units at a "high" price and reduces it until all units are soled.
- Sealed-bid – all bids are made simultaneously, and the item is sold to the highest bidder.

First-price / second-price

The price paid may be the highest bid or some other price:

- First-price – the bidder who submits the highest bid wins and pay a price equal to her bid.

- Second-prices – the bidder who submits the highest bid wins and pay a price equal to the second highest bid.

Variants: all-pay (lobbying), discriminatory, uniform, Vickrey (William Vickrey, Nobel Laureate 1996), and more.

Private-value / common-value

Bidders can be certain or uncertain about each other's valuation:

- In private-value auctions, valuations differ among bidders, and each bidder is certain of her own valuation and can be certain or uncertain of every other bidder's valuation.

- In common-value auctions, all bidders have the same valuation, but bidders do not know this value precisely and their estimates of it vary.

First-price auction (with perfect information)

To define the game precisely, denote by v_i the value that bidder i attaches to the object. If she obtains the object at price p then her payoff is $v_i - p$.

Assume that bidders' valuations are all different and all positive. Number the bidders 1 through n in such a way that

$$v_1 > v_2 > \cdots > v_n > 0.$$

Each bidder i submits a (sealed) bid b_i. If bidder i obtains the object, she receives a payoff $v_i - b_i$. Otherwise, her payoff is zero.

Tie-breaking – if two or more bidders are in a tie for the highest bid, the winner is the bidder with the highest valuation.

In summary, a first-price sealed-bid auction with perfect information is the following strategic game:

– Players: the n bidders.

– Actions: the set of possible bids b_i of each player i (nonnegative numbers).

– Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \bar{b} & \text{if} \quad b_i = \bar{b} \text{ and } v_i > v_j \text{ if } b_j = \bar{b} \\ 0 & \text{if} \quad b_i < \bar{b} \end{cases}$$

where $\bar{b}$ is the highest bid.

The set of Nash equilibria is the set of profiles $(b_1, ..., b_n)$ of bids with the following properties:

$$\begin{array}{ll} [1] & v_2 \leq b_1 \leq v_1 \\ [2] & b_j \leq b_1 \text{ \underline{for all} } j \neq 1 \\ [3] & b_j = b_1 \text{ \underline{for some} } j \neq 1 \end{array}$$

It is easy to verify that all these profiles are Nash equilibria. It is harder to show that there are no other equilibria. We can easily argue, however, that there is no equilibrium in which player 1 does not obtain the object.

$\Longrightarrow$ The first-price sealed-bid auction is socially efficient, but does not necessarily raise the most revenues.

Second-price auction (with perfect information)

A second-price sealed-bid auction with perfect information is the following strategic game:

- Players: the n bidders.

- Actions: the set of possible bids b_i of each player i (nonnegative numbers).

- Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \overline{b} & \text{if} \quad b_i > \overline{b} \text{ or } b_i = \overline{b} \text{ and } v_i > v_j \text{ if } b_j = \overline{b} \\ 0 & \text{if} \quad b_i < \overline{b} \end{cases}$$

  where $\overline{b}$ is the highest bid submitted by a player other than i.

First note that for any player i the bid $b_i = v_i$ is a (weakly) dominant action (a “truthful” bid), in contrast to the first-price auction.

The second-price auction has many equilibria, but the equilibrium $b_i = v_i$ for all i is distinguished by the fact that every player's action dominates all other actions.

Another equilibrium in which player $j \neq 1$ obtains the good is that in which

[1] $b_1 < v_j$ and $b_j > v_1$
[2] $b_i = 0$ for all $i \neq \{1, j\}$

Common-value auctions and the winner's curse

Suppose we all participate in a sealed-bid auction for a jar of coins. Once you have estimated the amount of money in the jar, what are your bidding strategies in first- and second-price auctions?

The winning bidder is likely to be the bidder with the largest positive error (the largest overestimate).

In this case, the winner has fallen prey to the so-called the winner's curse. Auctions where the winner's curse is significant are oil fields, spectrum auctions, pay per click, and more.

Evolutionary Game Theory

Evolutionary stability

A single population of players. Players interact with each other pair-wise and randomly matched.

Players are assigned modes of behavior (mutation). Utility measures each player's ability to survive.

ε of players consists of mutants taking action a while others take action a^*.

Evolutionary stable strategy (ESS)

Consider a two-player payoff symmetric game

$$G = \langle \{1, 2\}, (A, A), (u_1, u_2) \rangle$$

where

$$u_1(a_1, a_2) = u_2(a_2, a_1)$$

(players exchanging a_1 and a_2).

$a^* \in A$ is ESS if and only if for any $a \in A$, $a \neq a^*$ and $\varepsilon > 0$ sufficiently small

$$(1-\varepsilon)u(a^*, a^*) + \varepsilon u(a^*, a) > (1-\varepsilon)u(a, a^*) + \varepsilon u(a, a)$$

which is satisfied if and only if for any $a \neq a^*$ either

$$u(a^*, a^*) > u(a, a^*)$$

or

$$u(a^*, a^*) = u(a, a^*) \text{ and } u(a^*, a) > u(a, a)$$

Three results on ESS

[1] If a^* is an ESS then (a^*, a^*) is a NE.

Suppose not. Then, there exists a strategy $a \in A$ such that

$$u(a, a^*) > u(a^*, a^*).$$

But, for ε small enough

$$(1-\varepsilon)u(a^*, a^*) + \varepsilon u(a^*, a) < (1-\varepsilon)u(a, a^*) + \varepsilon u(a, a)$$

and thus a^* is not an ESS.

[2] If (a^*, a^*) is a strict NE $(u(a^*, a^*) > u(a, a^*)$ for all $a \in A)$ then a^* is an ESS.

Suppose a^* is not an ESS. Then either

$$u(a^*, a^*) \leq u(a, a^*)$$

or

$$u(a^*, a^*) = u(a, a^*) \text{ and } u(a^*, a) \leq u(a, a).$$

so (a^*, a^*) can be a NE but not a strict NE.

[3] The two-player two-action game

	a	a'
a	w, w	x, y
a'	y, x	z, z

has a strategy which is ESS.

If $w > y$ or $z > x$ then (a, a) or (a', a') are strict NE, and thus a or a' are ESS.

If $w < y$ and $z < x$ then there is a unique symmetric mixed strategy NE (α^*, α^*) where

$$\alpha^*(a) = (z - x)/(w - y + z - x)$$

and $u(\alpha^*, \alpha) > u(\alpha, \alpha)$ for any $\alpha \neq \alpha^*$.

Repeated games
(the prisoner's dilemma)

The basic idea – prisoner's dilemma

In the Prisoner's Dilemma

	C	D
C	3, 3	0, 4
D	4, 0	1, 1

No cooperation (D, D) is the unique NE since D strictly dominates C, but both players are better off when the outcome is (C, C).

When played repeatedly, cooperation (C, C) in every period is stable if

- each player believes that choosing D will end cooperation, and
- subsequent losses outweigh the immediate gain.

The socially desirable outcome (C, C) can be sustained if (and only if) players have long-term objectives.

In general, we can think that strategies are social norms, cooperation, threats and punishments where threats are carried out as punishments when the social norms require it.

Strategies

Grim trigger strategy

$$\boxed{\mathcal{C} : C} \underset{(\cdot, D)}{\longrightarrow} \boxed{\mathcal{D} : D}$$

Limited punishment

$$\dashrightarrow \boxed{\mathcal{P}_0 : C} \underset{(\cdot, D)}{\longrightarrow} \boxed{\mathcal{P}_1 : D} \underset{(\cdot, \cdot)}{\longrightarrow} \boxed{\mathcal{P}_2 : D} \underset{(\cdot, \cdot)}{\longrightarrow} \boxed{\mathcal{P}_3 : D} \underset{(\cdot, \cdot)}{\dashrightarrow}$$

Tit-for-tat

$$\dashrightarrow \boxed{\mathcal{C} : C} \underset{(\cdot, D)}{\longrightarrow} \boxed{\mathcal{D} : D} \underset{(\cdot, C)}{\dashrightarrow}$$

Payoffs

A player's preferences over an infinite stream $(\omega^1, \omega^2, \ldots)$ of payoffs are represented by the discounted sum

$$V = \sum_{t=1}^{\infty} \delta^{t-1} \omega^t,$$

where $0 < \delta < 1$.

The discounted sum of stream $(c, c, \ldots)$ is $\frac{C}{1-\delta}$, so a player is indifferent between the two streams if

$$c = (1-\delta)V.$$

Hence, we call $(1-\delta)V$ the discounted average of stream $(\omega^1, \omega^2, \ldots)$, which represent the same preferences.

To elucidate, let

$$S_T = c + \delta c + \delta^2 c + \cdots + \delta^T c$$

and note that

$$\delta S_T = \delta c + \delta^2 c + \delta^3 c + \cdots + \delta^{T+1} c$$

so that $S_T - \delta S_T = c - \delta^{T+1} c$ and thus

$$S_T = \frac{1 - \delta^{T+1}}{1 - \delta} c,$$

which equals $\dfrac{C}{1 - \delta}$ as $T \to \infty$.

Nash equilibria

Grim trigger strategy

$$(1-\delta)(3+\delta+\delta^2+\cdots) = (1-\delta)\left[3+\frac{\delta}{(1-\delta)}\right] = 3(1-\delta)+\delta$$

Thus, a player cannot increase her payoff by deviating if and only if

$$3(1-\delta)+\delta \leq 2,$$

or $\delta \geq 1/2$.

If $\delta \geq 1/2$, then the strategy pair in which each player's strategy is grim strategy is a Nash equilibrium which generates the outcome (C, C) in every period.

Limited punishment (k periods)

$$(1-\delta)(3+\delta+\delta^2+\cdots+\delta^k) = (1-\delta)\left[3+\delta\frac{(1-\delta^k)}{(1-\delta)}\right] = 3(1-\delta)+\delta(1-\delta^k)$$

Note that after deviating at period t a player should choose D from period $t+1$ through $t+k$.

Thus, a player cannot increase her payoff by deviating if and only if

$$3(1-\delta)+\delta(1-\delta^k) \leq 2(1-\delta^{k+1}).$$

Note that for $k=1$, then no $\delta < 1$ satisfies the inequality.

Tit-for-tat

A deviator's best-reply to tit-for-tat is to alternate between D and C or to always choose D, so tit-for tat is a best-reply to tit-for-tat if and only if

$$(1-\delta)(3+0+3\delta^2+0+\cdots) = (1-\delta)\frac{3}{1-\delta^2} = \frac{3}{1+\delta} \leq 2$$

and

$$(1-\delta)(3+\delta+\delta^2+\cdots) = (1-\delta)\left[3+\frac{\delta}{(1-\delta)}\right] = 3-2\delta \leq 2.$$

Both conditions yield $\delta \geq 1/2$.

Back to Decision Theory…

The fundamental tradeoffs in life

People's attitudes towards risk, time and other people enter every realm of (financial) decision-making:

$$\begin{array}{rcl} \text{risk} & \Longleftrightarrow & \text{return} \\ \text{today} & \Longleftrightarrow & \text{tomorrow} \\ \text{self} & \Longleftrightarrow & \text{others} \end{array}$$

Risk, time and social preferences are thus important inputs into any broader measure of welfare and enter virtually every field of economics.

Life is full of lotteries :-(

$$x := \begin{array}{cc} & \$A \\ \overset{p}{\nearrow} & \\ & \\ \underset{1-p}{\searrow} & \\ & \$B \end{array} \qquad y := \begin{array}{cc} & \$C \\ \overset{q}{\nearrow} & \\ \xrightarrow{q'} & \$D \\ \underset{1-q-q'}{\searrow} & \\ & \$E \end{array}$$

A risky lottery (left) and an ambiguous lottery (right)

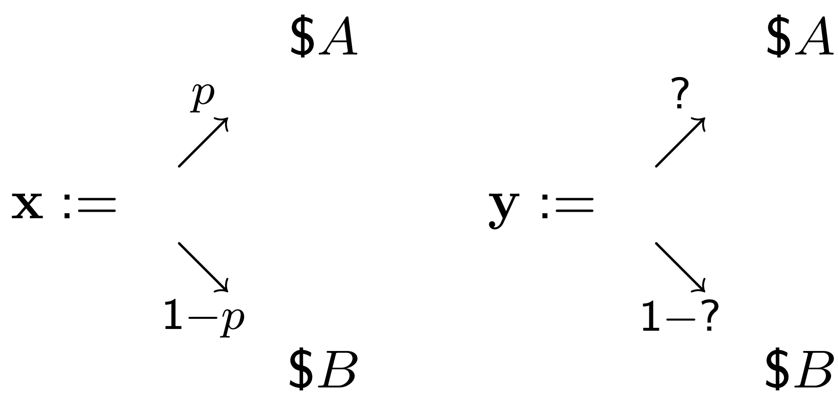

A compounded lottery

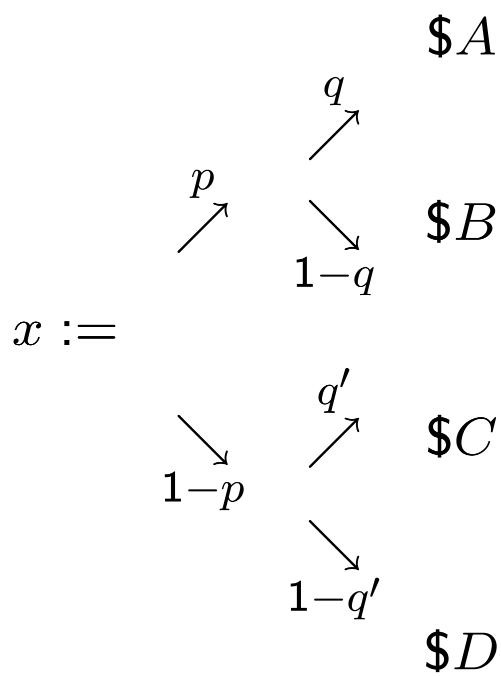

The fundamental assumptions (axioms) about (risk) preferences

All theories (EU and non-EU) begin with three assumptions about preferences:

Completeness

For any pair of lotteries or gambles (outcomes and probabilities) x and y

$$x \succsim y \text{ or } y \succsim x.$$

The fundamental assumptions (axioms) about (risk) preferences

All theories (EU and non-EU) begin with three assumptions about preferences:

Transitivity

For any three lotteries x, y, z

$$\text{if } x \succsim y \text{ and } y \succsim z \text{ then } x \succsim z.$$

The fundamental assumptions (axioms) about preferences

All theories (EU and non-EU) begin with three assumptions about preferences:

Monotonicity (with respect to first-order stochastic dominance)

For any pair of lotteries x and y with resulting payoff distributions F_x and F_y

$$\text{if } F_x \geq F_y \text{ then } x \succsim y.$$

$\Rightarrow$ The preferences can be represented, or summarized, by a well-behaved (increasing) utility function.

The hunt for a descriptive theory of choice under risk (Starmer, 2000)

The 'standard' model of decisions under risk is based on von Neumann and Morgenstern Expected Utility (EU):

Independence

For any three lotteries x, y, z and $0 < \alpha < 1$

$$\text{if } x \succ y \text{ then } \alpha x + (1-\alpha)z \succ \alpha y + (1-\alpha)z.$$

$\Rightarrow$ Empirical violations of independence generated the development of various theoretical alternatives, and the investigation of these theories has led to new empirical regularities, and so on...

Independence

$$x := \begin{array}{c} \overset{p_x}{\nearrow} \\ \\ \underset{1-p_x}{\searrow} \end{array} \begin{array}{c} \$A \\ \\ \$B \end{array} \quad \succ \quad y := \begin{array}{c} \overset{p_y}{\nearrow} \\ \\ \underset{1-p_y}{\searrow} \end{array} \begin{array}{c} \$C \\ \\ \$D \end{array}$$

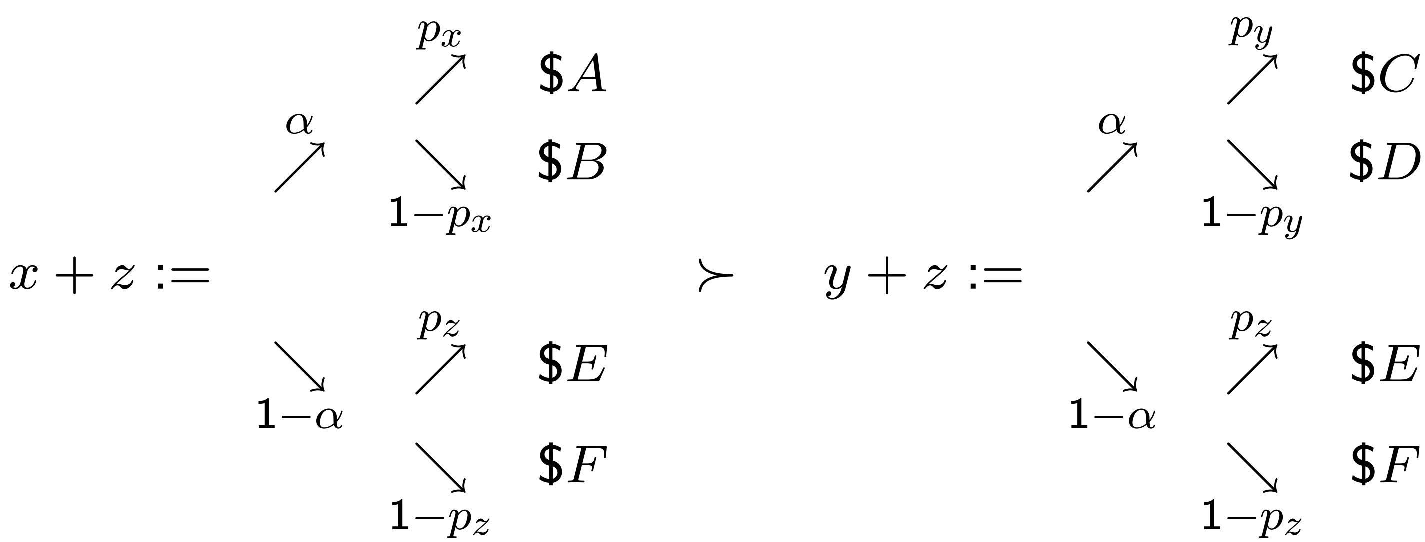
$x + z :=$
α
$1-\alpha$
p_x
$1-p_x$
p_z
$1-p_z$
$\$A$
$\$B$
$\$E$
$\$F$
$\succ$
$y + z :=$
α
$1-\alpha$
p_y
$1-p_y$
p_z
$1-p_z$
$\$C$
$\$D$
$\$E$
$\$F$

The Allais paradox (1953)

Allais (1953) I

Choose between the two lotteries:

$$x := \begin{array}{c} \overset{.33}{\nearrow} \\ \xrightarrow{.66} \\ \underset{.01}{\searrow} \end{array} \begin{array}{c} \$25,000 \\ \$24,000 \\ \$0 \end{array} \qquad y := \xrightarrow{1} \$24,000$$

Allais (1953) II

Choose between the two lotteries:

$$x' := \begin{array}{c} .33 \\ \nearrow \\ \\ \searrow \\ .67 \end{array} \begin{array}{c} \$25,000 \\ \\ \$0 \end{array} \qquad y' := \begin{array}{c} .34 \\ \nearrow \\ \\ \searrow \\ .66 \end{array} \begin{array}{c} \$24,000 \\ \\ \$0 \end{array}$$

The (Marschak-Machina) probability triangle

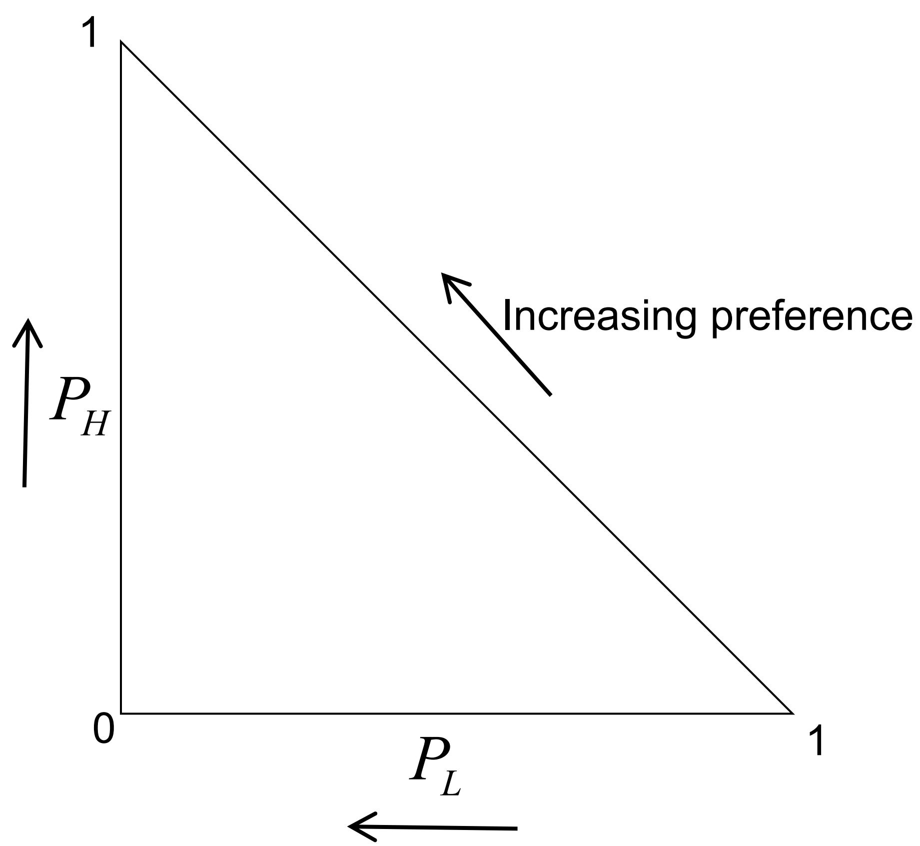

Consider three monetary payouts H, *M*, and *L* where *H>M>L*

An indifference map of a loss-neutral (expected utility) individual

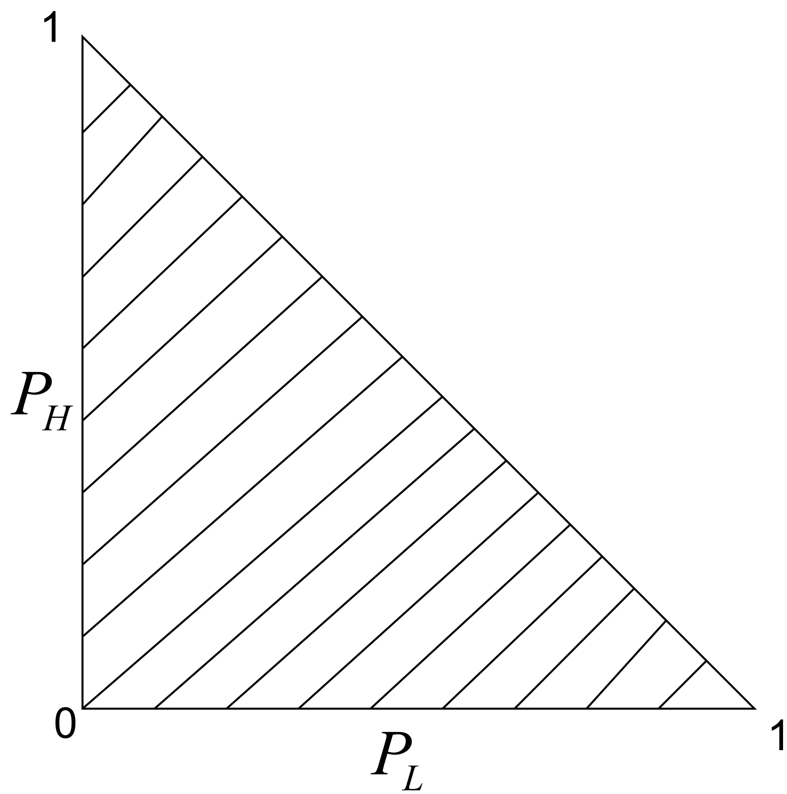

Expected Utility Theory (EUT) requires that indifference lines are parallel

A test of Expected Utility Theory (EUT)

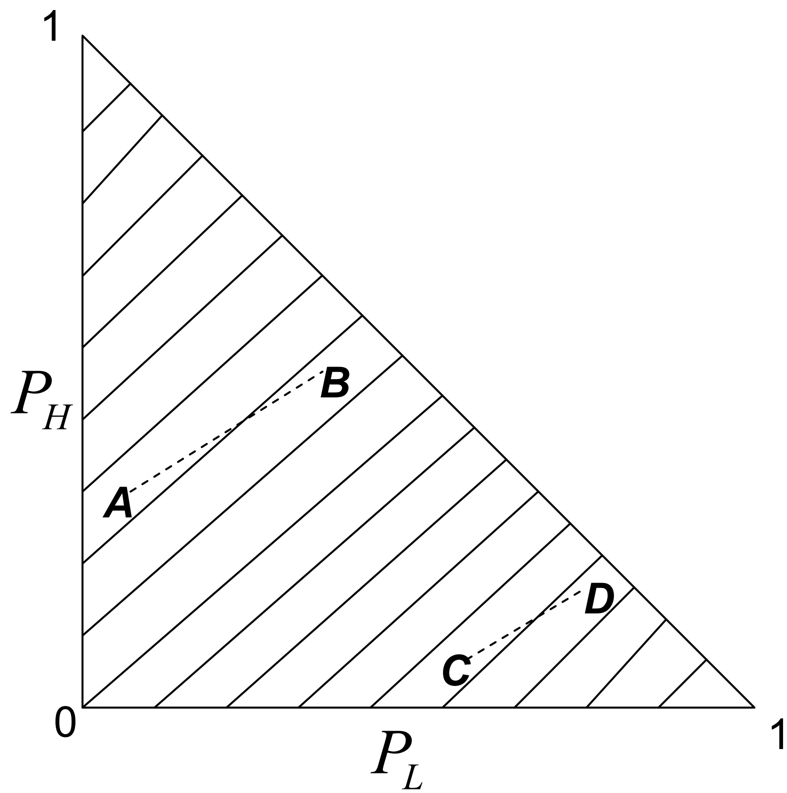

EUT requires that indifference lines are parallel so one must choose either ***A*** and ***C***, or ***B*** and ***D***.

The decision problem

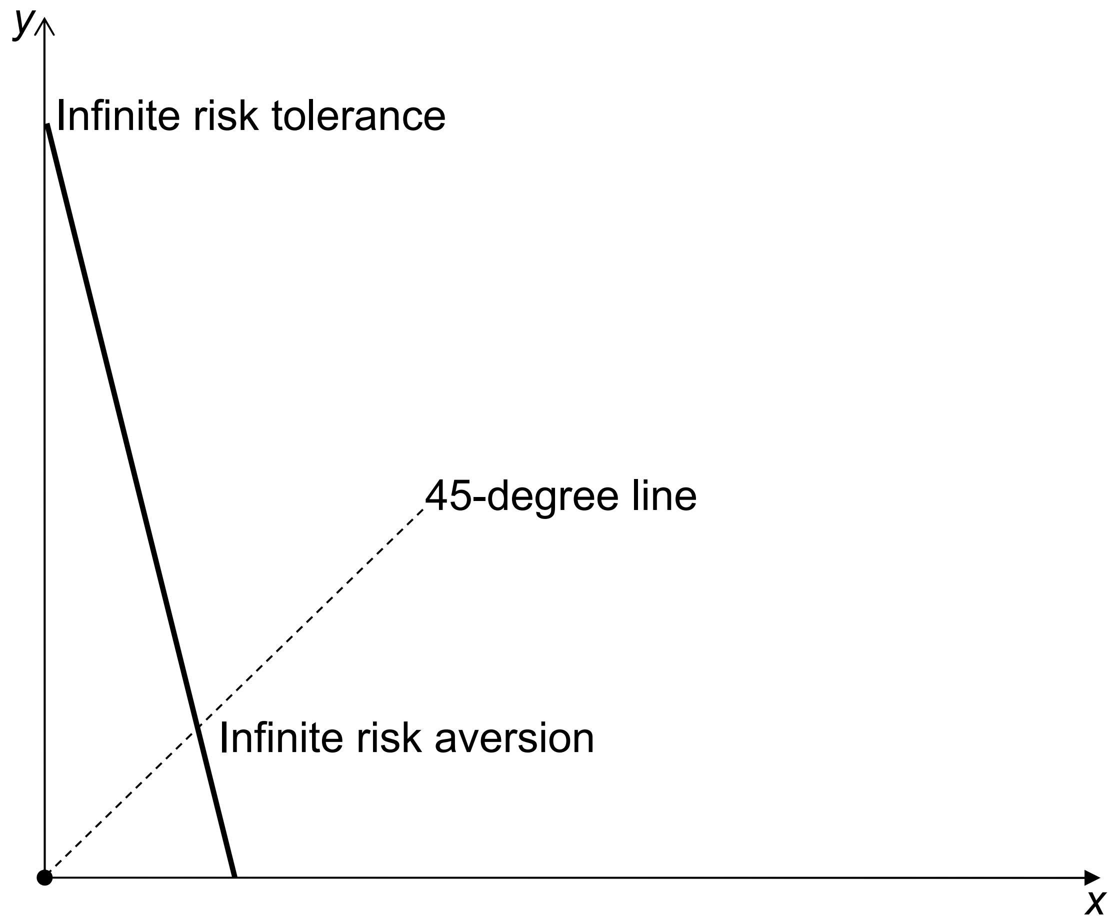

Individual-level data

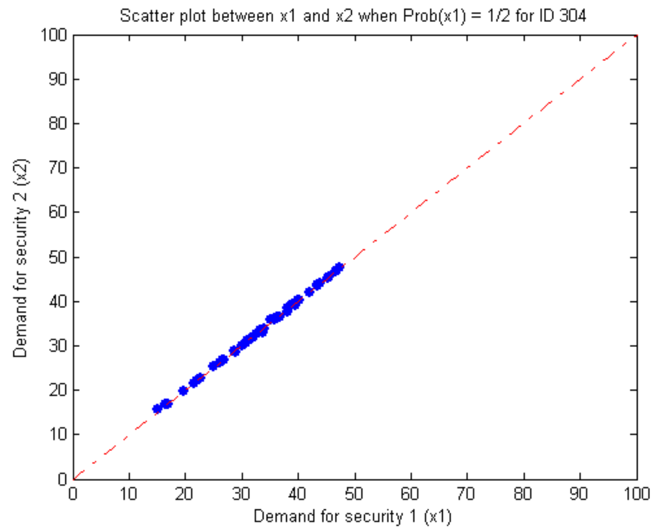

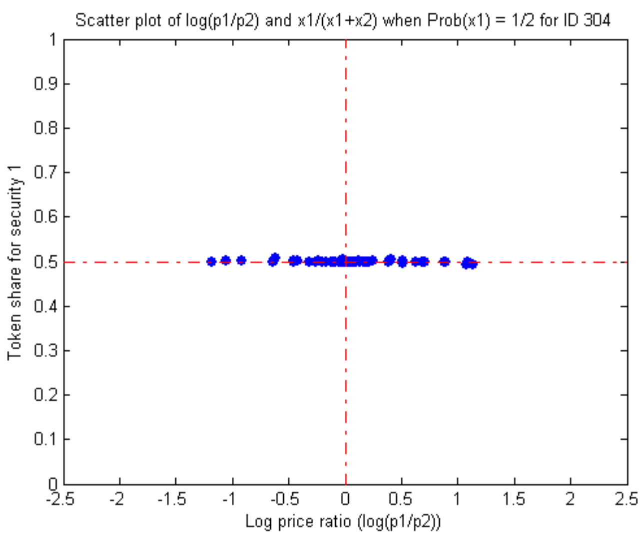

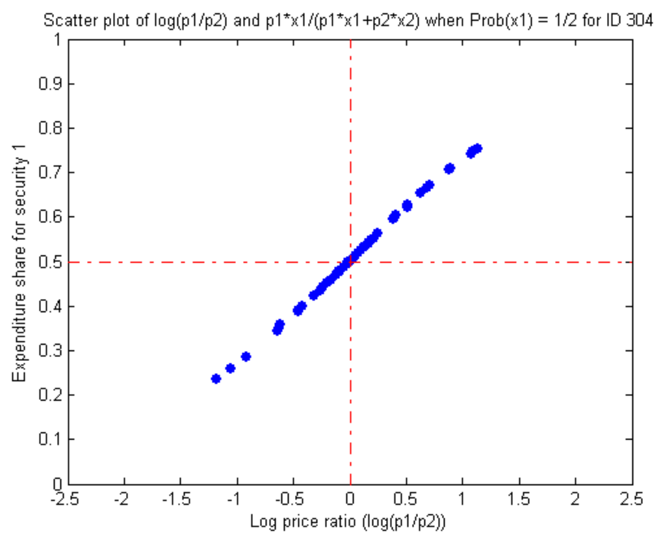

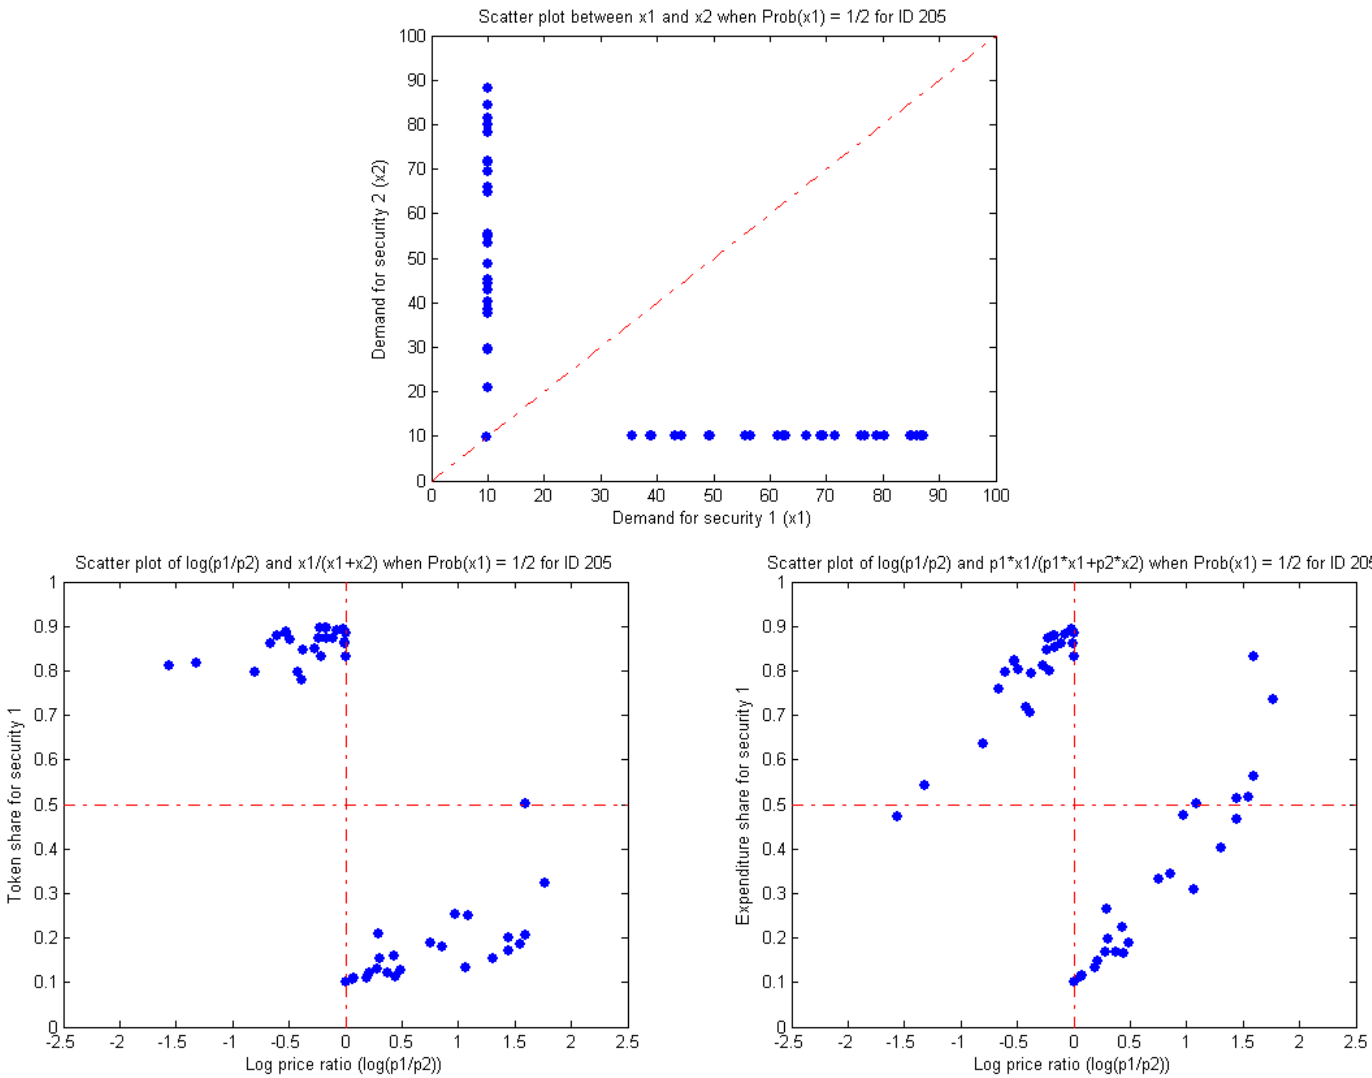

Scatter plot between x1 and x2 when Prob(x1) = 1/2 for ID 205
Demand for security 2 (x2)
Demand for security 1 (x1)
Scatter plot of log(p1/p2) and x1/(x1+x2) when Prob(x1) = 1/2 for ID 205
Token share for security 1
Log price ratio (log(p1/p2))
Scatter plot of log(p1/p2) and p1*x1/(p1*x1+p2*x2) when Prob(x1) = 1/2 for ID 205
Expenditure share for security 1
Log price ratio (log(p1/p2))

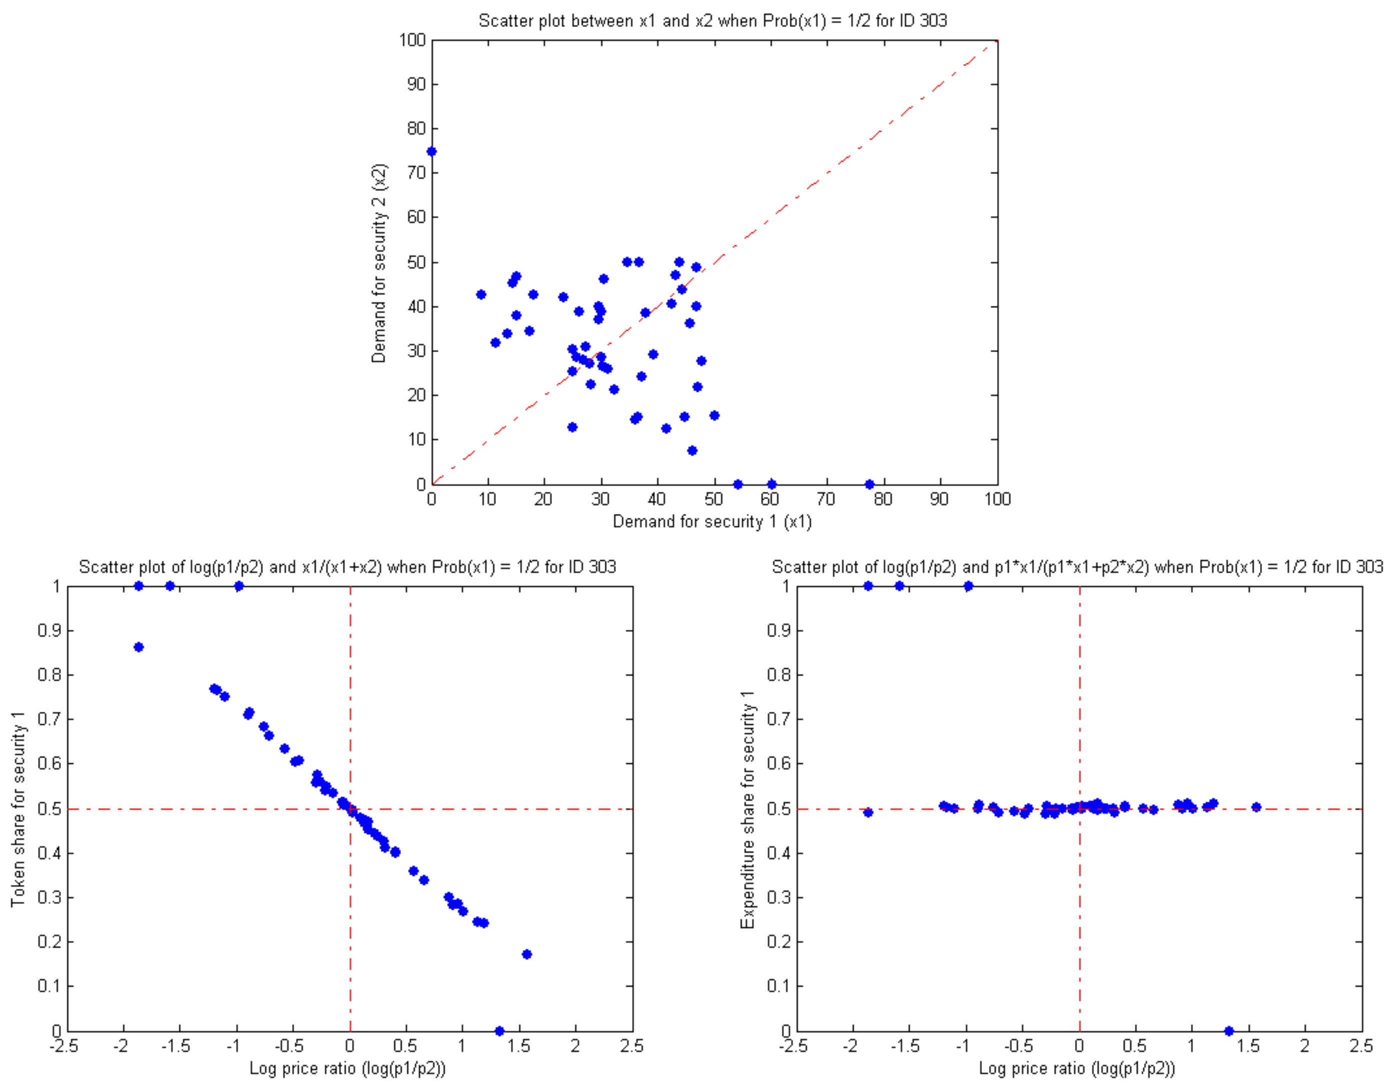
Scatter plot between x1 and x2 when Prob(x1) = 1/2 for ID 303
Demand for security 2 (x2)
Demand for security 1 (x1)
Scatter plot of log(p1/p2) and x1/(x1+x2) when Prob(x1) = 1/2 for ID 303
Token share for security 1
Log price ratio (log(p1/p2))
Scatter plot of log(p1/p2) and p1*x1/(p1*x1+p2*x2) when Prob(x1) = 1/2 for ID 303
Expenditure share for security 1
Log price ratio (log(p1/p2))

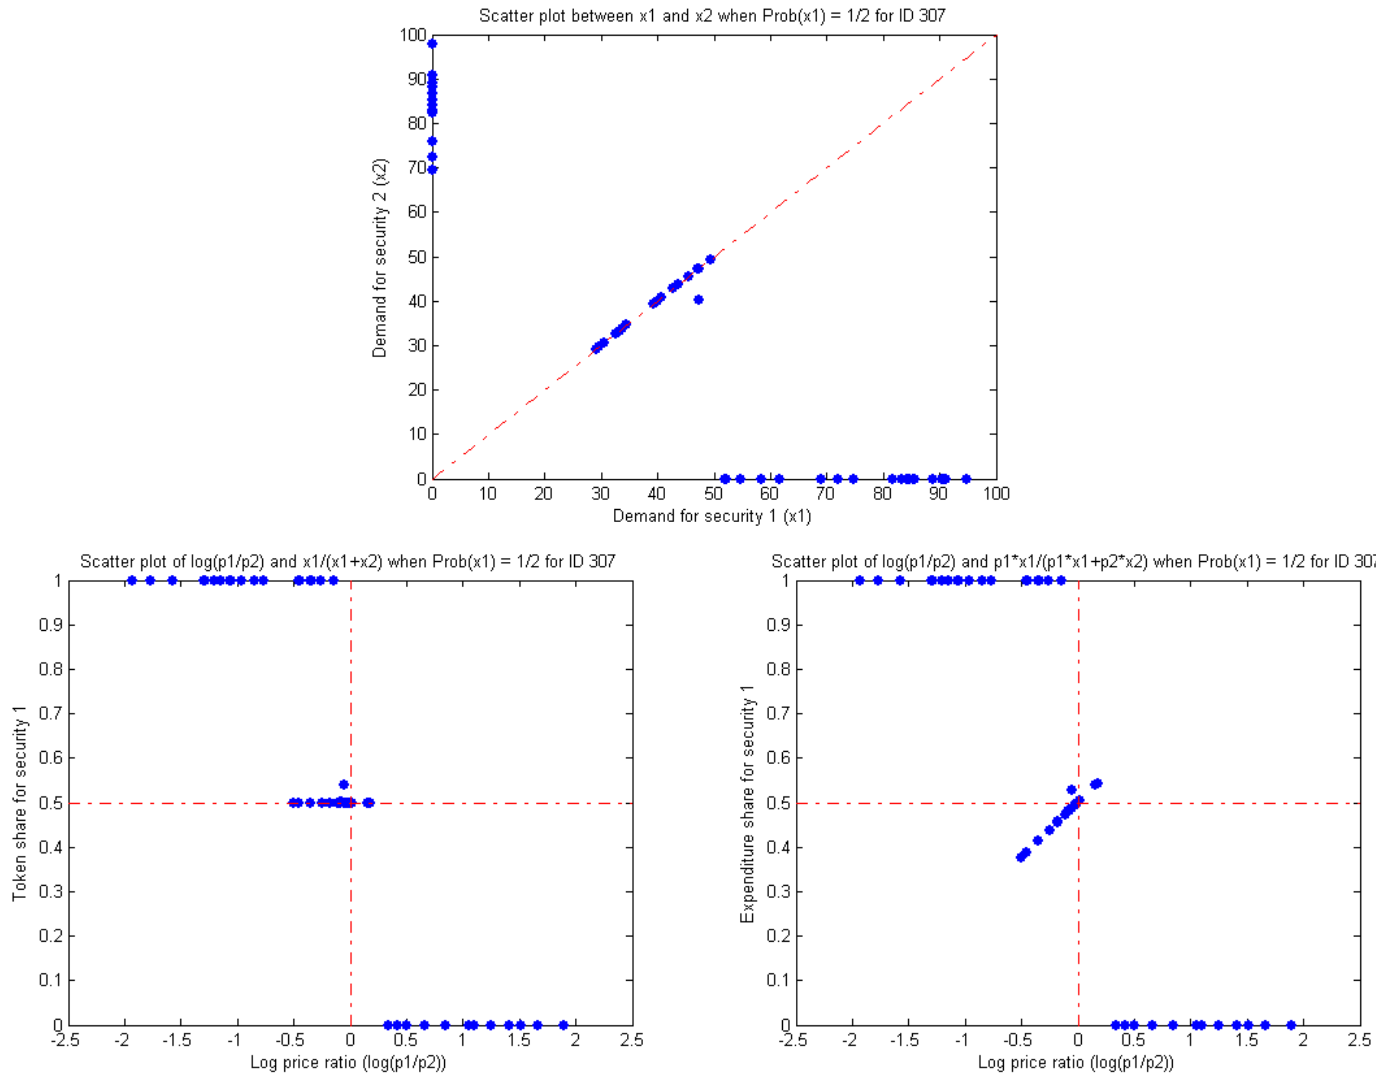
Scatter plot between x1 and x2 when Prob(x1) = 1/2 for ID 307
Demand for security 2 (x2)
Demand for security 1 (x1)
0
10
20
30
40
50
60
70
80
90
100
Scatter plot of log(p1/p2) and x1/(x1+x2) when Prob(x1) = 1/2 for ID 307
Token share for security 1
Log price ratio (log(p1/p2))
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
-2.5
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2
2.5
Scatter plot of log(p1/p2) and p1*x1/(p1*x1+p2*x2) when Prob(x1) = 1/2 for ID 307
Expenditure share for security 1
Log price ratio (log(p1/p2))

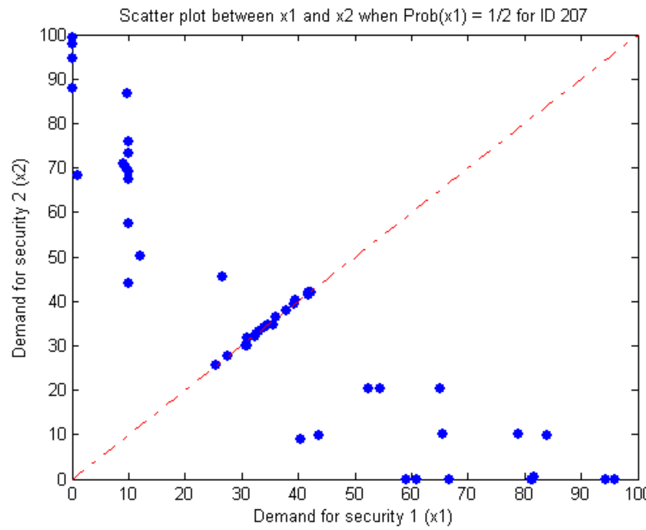
Scatter plot between x1 and x2 when Prob(x1) = 1/2 for ID 207
Demand for security 2 (x2)
Demand for security 1 (x1)

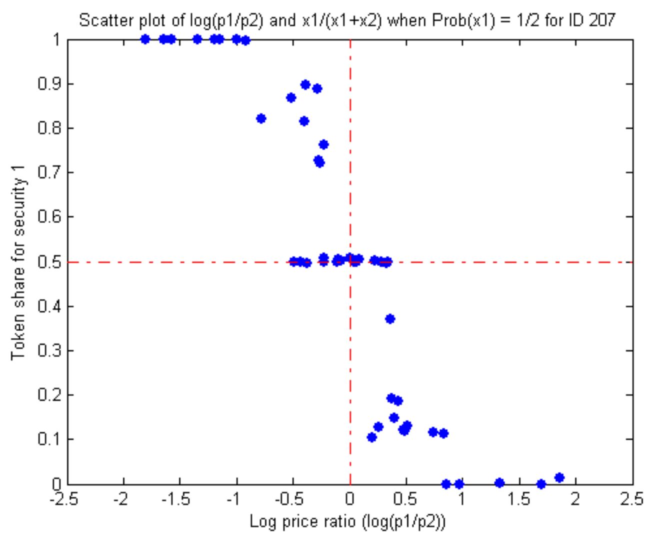
Scatter plot of log(p1/p2) and x1/(x1+x2) when Prob(x1) = 1/2 for ID 207
Token share for security 1
Log price ratio (log(p1/p2))

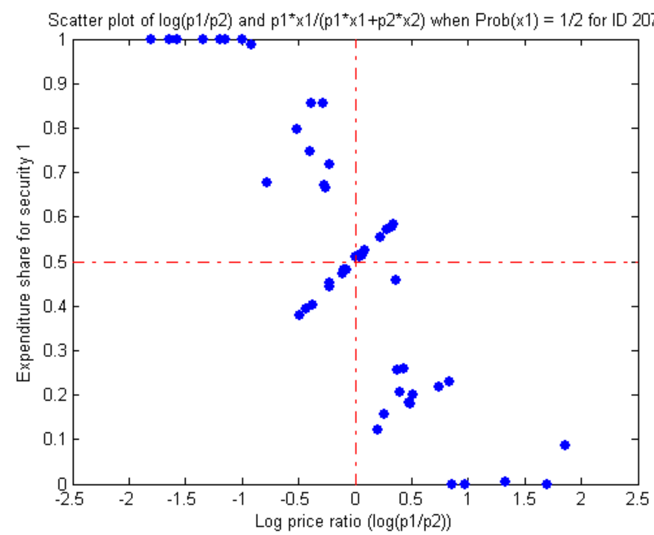
Scatter plot of log(p1/p2) and p1*x1/(p1*x1+p2*x2) when Prob(x1) = 1/2 for ID 207
Expenditure share for security 1
Log price ratio (log(p1/p2))

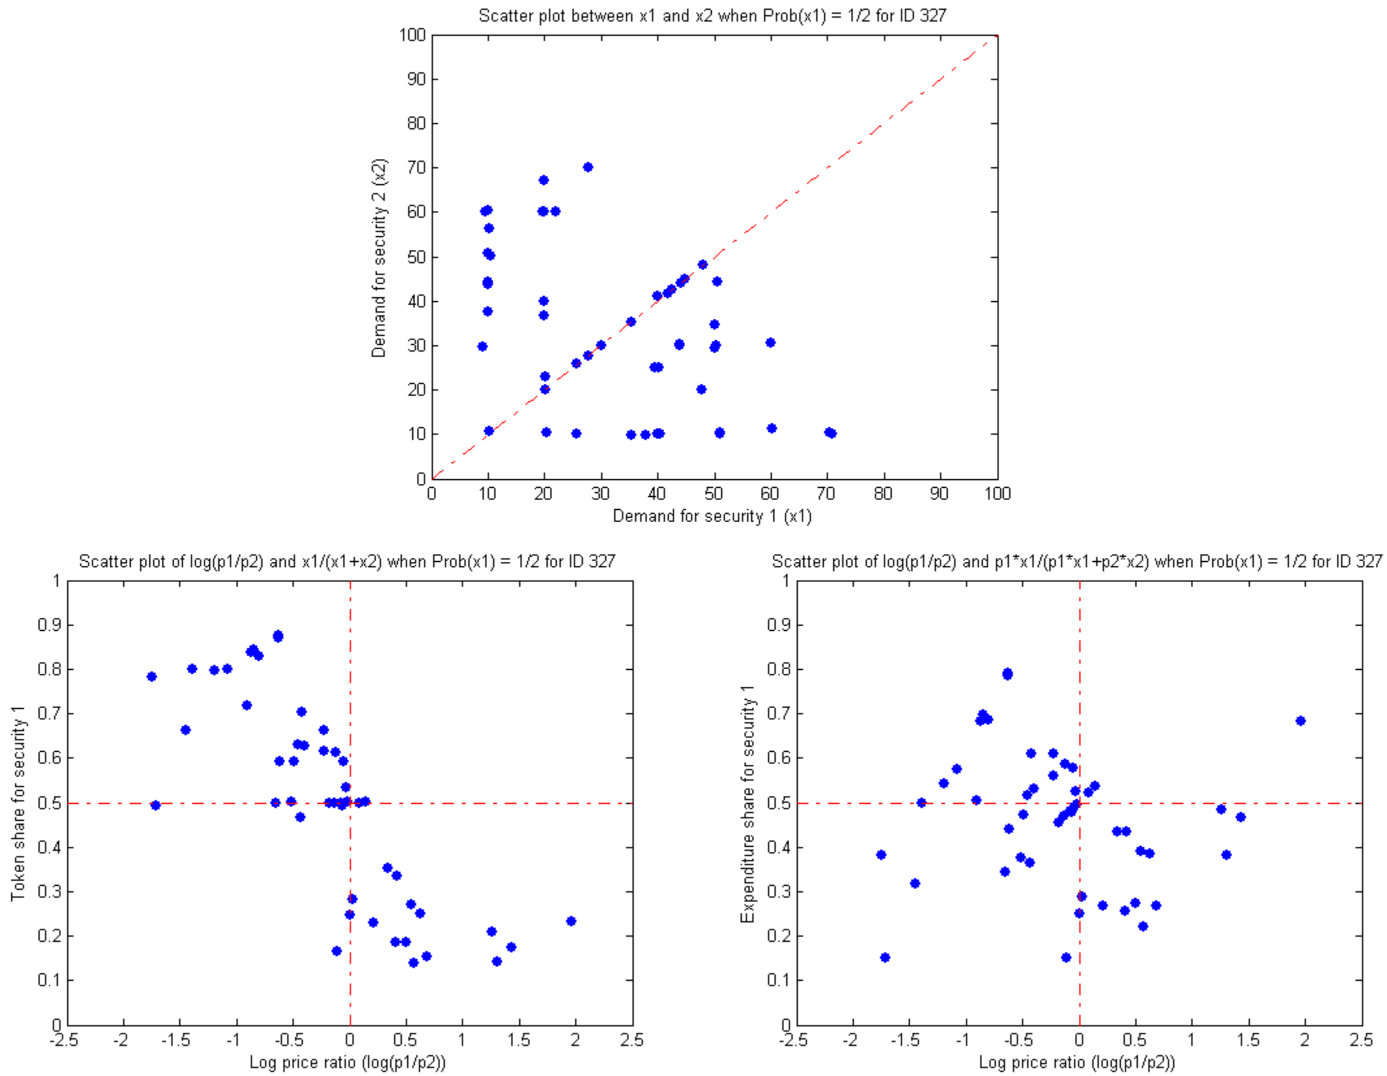
Scatter plot between x1 and x2 when Prob(x1) = 1/2 for ID 327
Demand for security 2 (x2)
Demand for security 1 (x1)
Scatter plot of log(p1/p2) and x1/(x1+x2) when Prob(x1) = 1/2 for ID 327
Token share for security 1
Log price ratio (log(p1/p2))
Scatter plot of log(p1/p2) and p1*x1/(p1*x1+p2*x2) when Prob(x1) = 1/2 for ID 327
Expenditure share for security 1
Log price ratio (log(p1/p2))

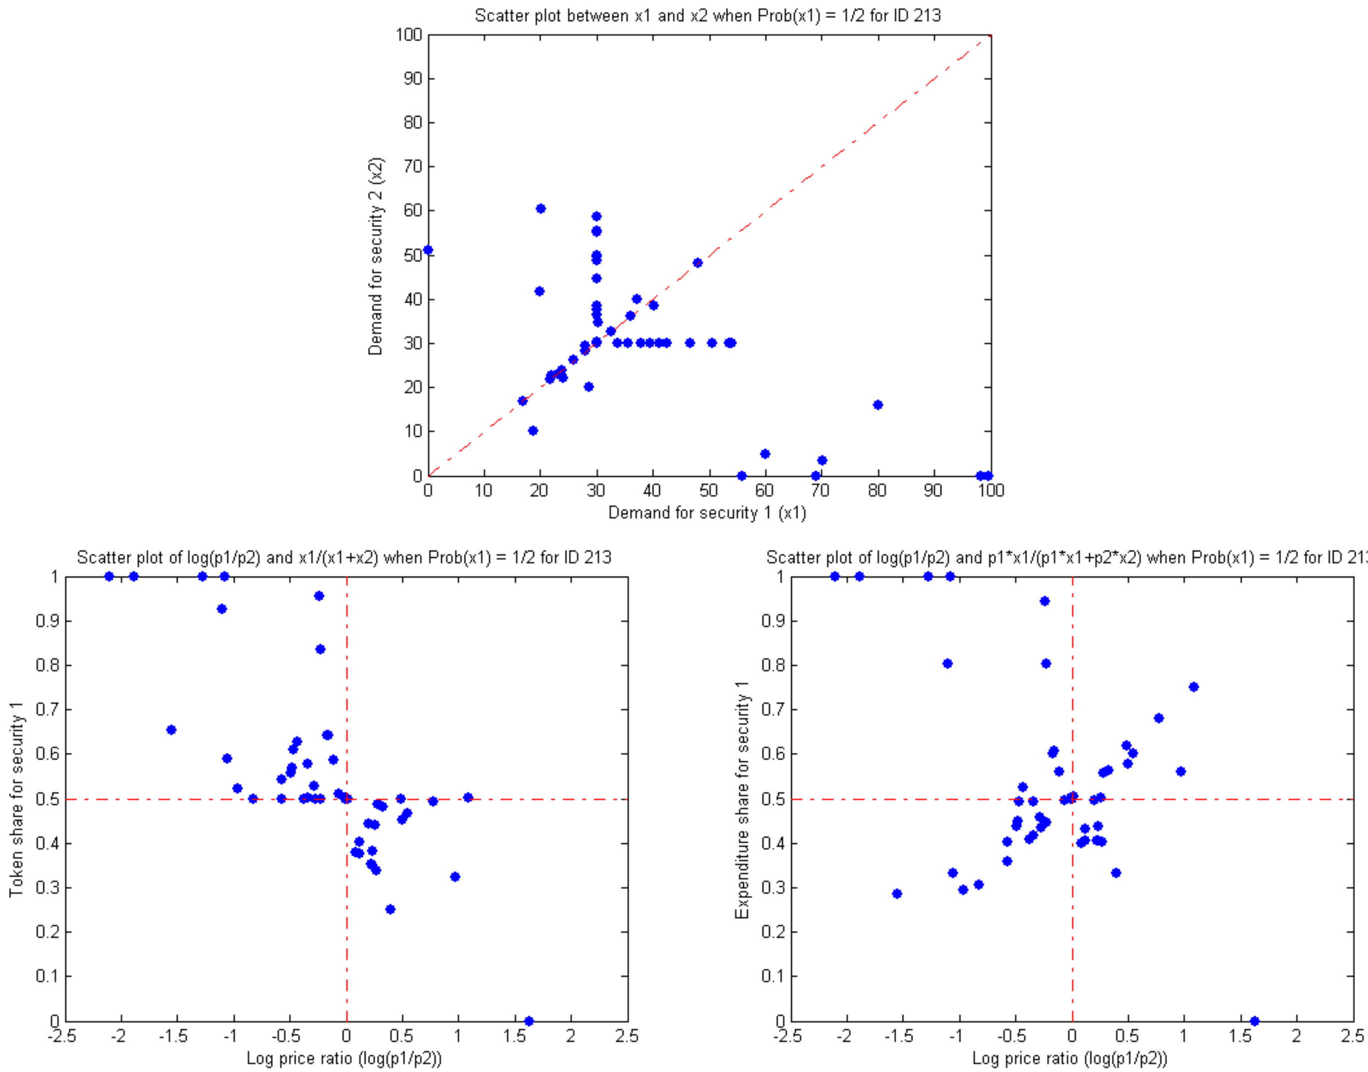

Scatter plot between x1 and x2 when Prob(x1) = 1/2 for ID 213
Demand for security 2 (x2)
Demand for security 1 (x1)
Scatter plot of log(p1/p2) and x1/(x1+x2) when Prob(x1) = 1/2 for ID 213
Token share for security 1
Log price ratio (log(p1/p2))
Scatter plot of log(p1/p2) and p1*x1/(p1*x1+p2*x2) when Prob(x1) = 1/2 for ID 213
Expenditure share for security 1
Log price ratio (log(p1/p2))

The construction of the CCEI for a simple violation of GARP

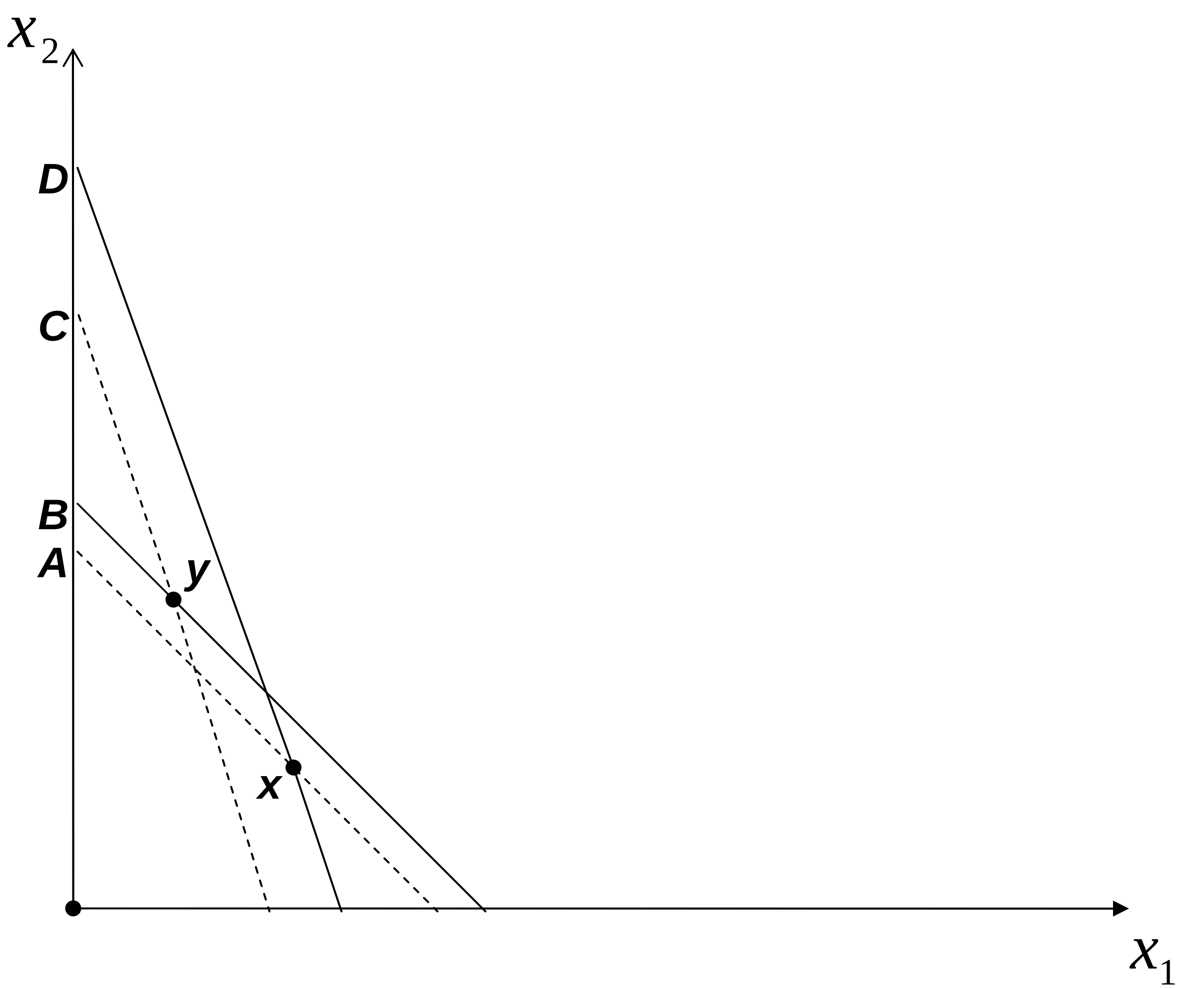

The agent is ‘wasting' as much as *A/B*<*C/D* of his income by making inefficient choices.

- e^* – maximizing any utility function (GARP).
- $e^{**} \leq e^*$ – maximizing a monotonic utility function (GARP+FOSD).
- $e^{***} \leq e^{**}$ – maximizing an expected utility function (GARP+FOSD+EU).

$\Rightarrow$ For all non-EU theories, which number well into double figures (Starmer, 2000)

$$e^{***} < e^{**} = e^* \ = 1.$$

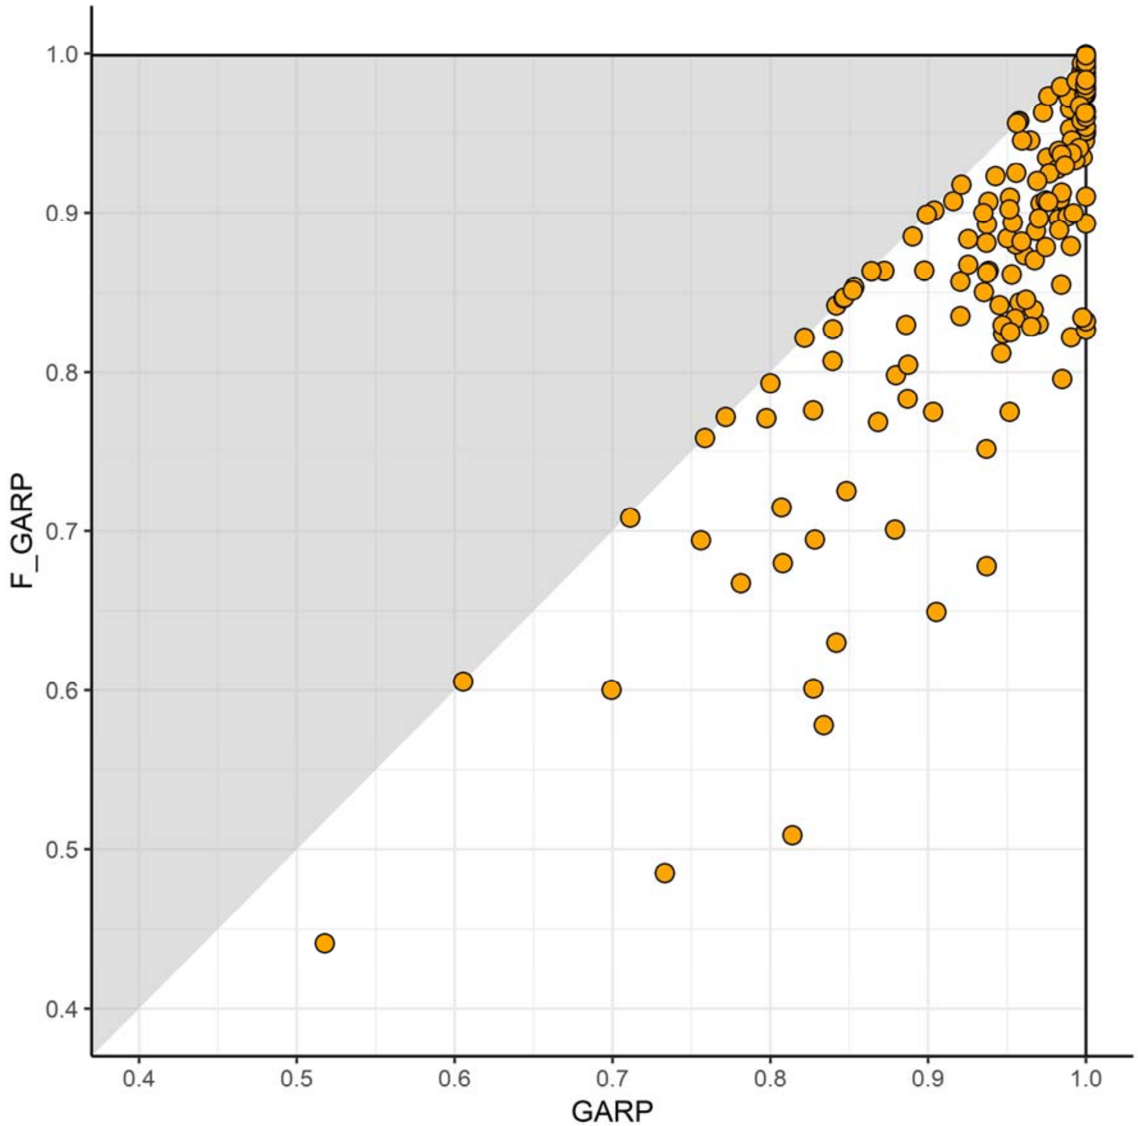

1.0
0.9
0.8
0.7
0.6
0.5
0.4
F_GARP
0.4
0.5
0.6
0.7
0.8
0.9
1.0
GARP

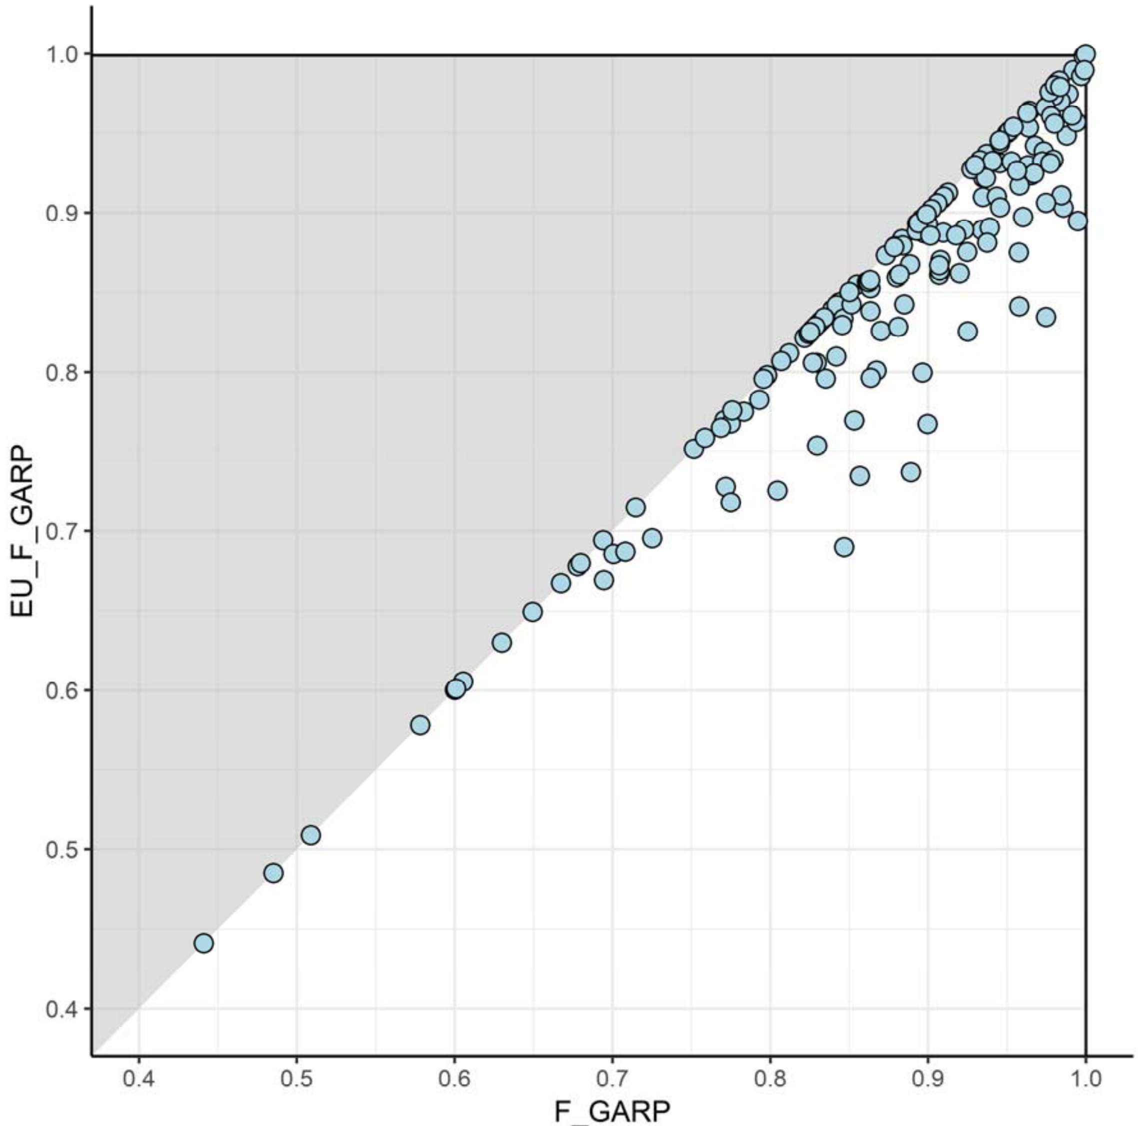
EU_F_GARP
F_GARP
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.4
0.5
0.6
0.7
0.8
0.9
1.0

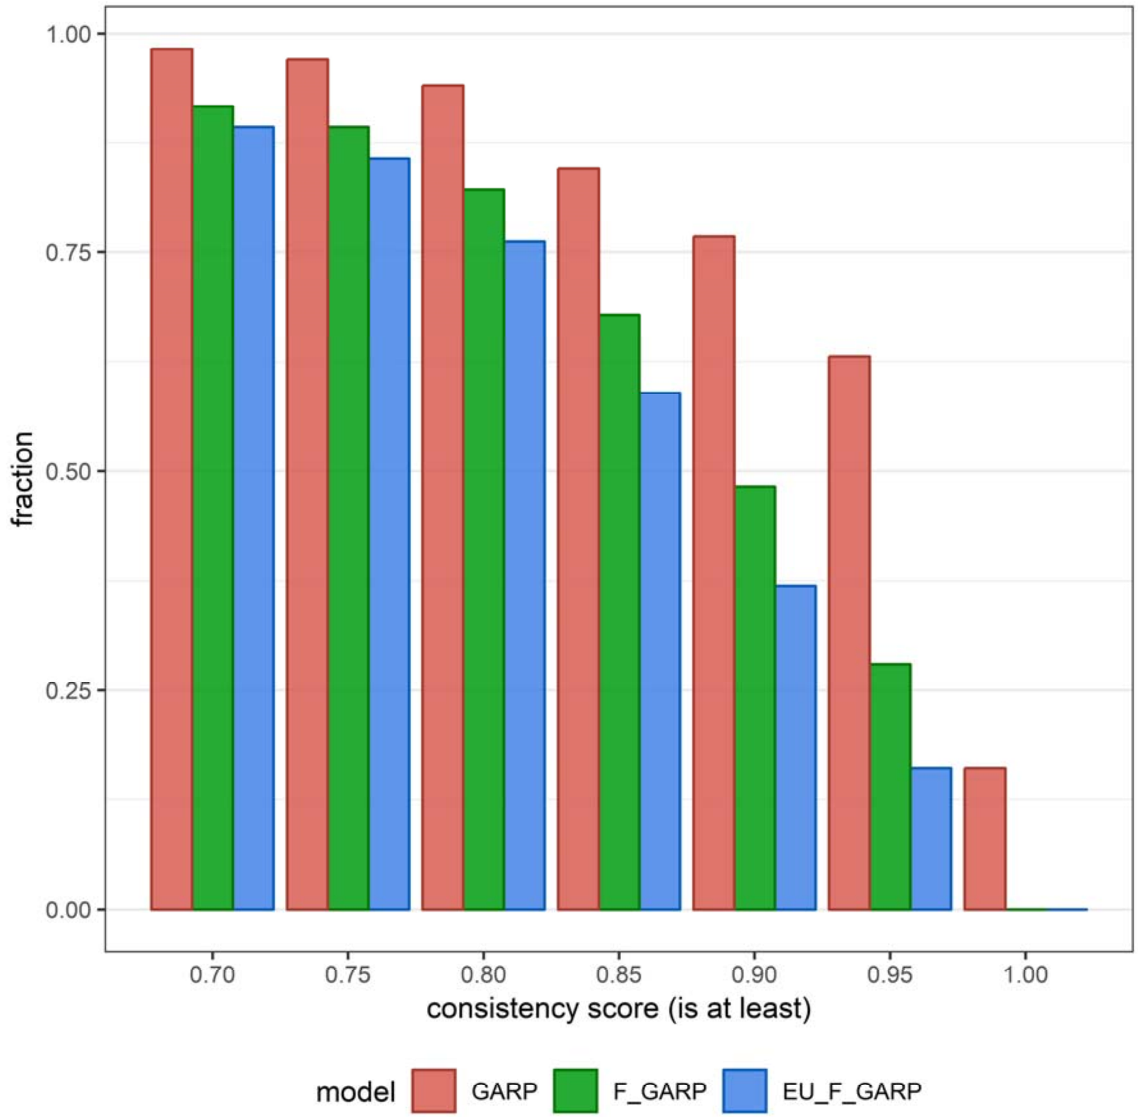

fraction
1.00
0.75
0.50
0.25
0.00
0.70
0.75
0.80
0.85
0.90
0.95
1.00
consistency score (is at least)
model
GARP
F_GARP
EU_F_GARP