K-level reasoning in beliefs

Preliminary draft. Please do not circulate

Marco Serena[*][†]

March 7, 2017

Abstract

In a Bayes-Nash Equilibrium of a private information game players engage in an iterative beliefs formation process of the form "I believe that you believe that I believe...", and so on, ad infinitum. However, in reality beliefs might extend only a few steps. We propose a non-equilibrium concept in which a player is Lk (i.e., her depth of reasoning is k) if she correctly forms up to k^{th}-order belief. Thus, an L0 does not process the reciprocal belief formation process of the game *at all*, in that she simplistically believes that her rival is of her same type.

We propose a simple game to test our level-k concept, and we show its prediction in standard games. In some games there is a sharp discontinuity between the infinite unraveling of reciprocal beliefs formation process and the Bayes-Nash Equilibrium.

Keywords: k-level reasoning, private information.

C72, C91, D83

1 Introduction

Consider a game with two players, Anne and Betty, each having private information about her own *type* (i.e., some payoff-relevant parameters), with types distributed according to the distribution F. The conventional game theoretic approach dictates that Anne's strategy depends on her expectation on Betty's type (first-order expectation), but also on her expectation on Betty's first-order expectation about Anne's type (second-order expectation), and so on ad infinitum. The infinite iteration of such reciprocal expectations which lies behind the definition of a Bayesian Nash equilibrium is inherently difficult to be fully accounted in a player's mental

[*]Max Planck Institute for Tax Law and Public Finance. *Email*: marco.serena@tax.mpg.de

[†]I would especially like to thank Tilman Klumpp for many substantial comments. I am also grateful to Vincent Crawford for useful suggestions. I would like to thank participants at the V Informal Research Meeting at the Max Planck Institute for Tax Law and Public Finance. All errors are my own.

process.[1] Based on this, we propose a model to capture the idea that players only engage in a *limited* number of iterations of reciprocal expectations. In particular, we define the "depth of reasoning" k of a player as the number of iterated expectations that the player computes; if Anne's depth of reasoning is $k = 1$ — shortly, Anne is an "L1" —, then Anne 1) understands that Betty's type follows F, but 2) does not understand that Betty forms beliefs on Anne's types according to F, and rather believes that Betty acts as an L0. Similarly, if instead Anne is L2, she 1) understands that Betty's type follows F, 2) believes that Betty understands that Anne's type follows F, but 3) believes that Betty believes that Anne acts as an L0. In other words, an Lk-player forms *correctly* up to the k^{th}-order beliefs, and when her belief formation process stops after k steps, the L0 behavior is assumed. Thus, it is crucial to carefully specify as anchoring point the expectation of an L0 — as is standard in models of level-k reasoning. Since an Lk iterates "I understand that my rival's type follows F" exactly k times, it is natural for an L0 to completely neglect the prior F, and simply believe that her rival is of her *same* type. More formally, an L0 of a certain type $\bar{\theta}$ (drawn from F) acts as if playing against another $\bar{\theta}$-player with certainty.[2] Thus, in a compact logic sentence, we propose the following non-equilibrium concept:

An Lk-player {believes that her rival's type ***does*** *follow F, and that her rival in turn}^{k} believes that her rival is of type which* ***does not*** *follow F, and rather is of her same type.*[3]

The closest literature is the one of level-k reasoning (e.g. Stahl and Wilson, 1994; Nagel, 1995; Costa-Gomes, Crawford,and Broseta, 2001), which mostly focused on games with complete information. Two key-ingredients characterize a level-k model: (i) the role of k in affecting actions, and (ii) the anchoring strategy of an L0. In the standard complete information setting an Lk iterates k times the best reply,[4] and the usual assumption on the anchoring strategy of an L0 is *behavioral*; an L0 randomly chooses among all the possible actions, thus neglecting payoffs and best reply functions.[5] In a private information setting, the best reply itself requires some

[1] Harsanyi (1967) provided groundbreaking tools to simplify the analytical tractability of models under incomplete information.

[2] Throughout the paper we assume type-symmetry; that is, identical types choose identical strategies. This assumption avoids multiplicity issues in the strategy of L0-players in some particular games. Also, this assumption is natural in the light of the fact that an L0 believes that her rival is identical to herself.

One could alternatively consider that an L0 believes that her rival chooses the same *action*, rather than is of the same type, as she does. Then L0 behavior can be justified by means of the concept of self-similarity, as introduced by Rubinstein and Salant (2016): self-similarity means that a player who chooses some action X tends to believe, to a greater extent than a player who chooses a different action, that other players will also choose action X.

[3] The part in the curly brackets is to be repeated exactly k times, so that an Lk computes correctly up to k^{th}-order of beliefs. Throughout the paper, the extension from 2-player to n-player is often trivial, and thus it will be omitted unless differently stated.

[4] Despite the vast majority of the level-k models assume that an Lk best replies to all other being L$(k-1)$, a remarkable alternative is that of cognitive hierarchy theory (Camerer, Ho and Chong, 2004) where an Lk best replies to a mix of L$(k-1)$,L$(k-2)$,...L0-players.

[5] An alternative L0, proposed by Crawford and Iriberri (2007) in private-value auctions, is that the bidders simply bid the value suggested by their own private information. We model L0 as

assumption on beliefs (I react differently if I am more likely to be up against a high type), and thus we model the depth of reasoning k as the order of beliefs up to which an Lk *correctly* reasons. This way, k captures the player's depth of understanding of the role of information in the game, which is embodied in the distribution F. In line with this, an L0 neglects the role of F already when computing her first-order beliefs, which are thus *wrong*. In particular, we argue that a natural belief for an "instinctive" L0 who does not understand the role of F is to have the most simplistic beliefs possible; "my rival is like me".[6] In Subsection 1.2 we argue that these simplistic beliefs of an L0 can be interpreted as a cognitive bias, which is supported by a well-established literature in psychology.

In what follows, we introduce a game suitable to capture the "depth of reasoning" of players in private information games. We present here its simplest version, for the sake of the argument, and we generalize and formally analyze it in Section 2.

1.1 The Up-or-Down game

Each of two identical decks contain cards from 1 to M.[7] Each player draws and privately observes a random card, of value $\theta_i \in \{1, .., M\}$, from her deck. Players simultaneously put on the table their cards either face Down (action D) or face Up (action U). The payoff matrix for player i is as follows,

		j's action	
		D	U
i's action	D	0	1
	U	ε	$\varepsilon + \alpha_i$

where $\varepsilon > 0$ is arbitrarily small, and

playing against an identical type with certainty, and thus in private-value first-price auctions our L0 who privately observes that her valuation is V believes that her rival's valuation is also V and that there is complete information. This yields a common-value complete-information first-price auction, where it is known that in equilibrium bidders simply bid their own valuation. Thus, in this sense, our assumption on L0 rationalizes Crawford and Iriberri's behavioral L0 assumption. See Subsection 4.5 for further discussions.

[6] As mentioned, the main reason behind our assumption on L0 is that we want the depth of reasoning k to capture the depth of understanding of the role of information in the game. However, there are other advantages of it as opposed to the standard randomization over the action space. In fact, our L0 behavior, as opposed to the standard randomization,

i) is not affected by topologic features of the strategy space (e.g., the L0's action is unambiguously defined in our setting when the game has an unbounded strategy space)

ii) is not affected by non-linear qualitatively-invariant changes in the action variables (e.g., is the players' action translate non-linearly into an impact which affects payoffs, our L0 behavior can be equivalently specified in action or impact)

iii) is likely to pin down a pure-strategy equilibrium (an exception is the all-pay auction, as discussed below), which are easier to identify in the lab than mixed strategies

iv) allows to identify, when moving from L0 to L1, the impact of the formation of first-order beliefs on actions, rather than the role of belief formation and best reply computation if we assumed that L0 randomize uniformly.

v) it does not presume on the game specifics, and thus it is in some sense more portable across games than behavioral assumptions such as the most "attractive" action in a specific game.

[7] Formally, in this introductory example, we need $M > 4$, and the reason (rather technical) will be clear in Section 2 and called *(A3)*. See also "**Remark on** M" in Section 2.

$$\alpha_i = \begin{cases} 4 & \text{if } \theta_i > \theta_j \\ 2 & \text{if } \theta_i = \theta_j \\ 0 & \text{if } \theta_i < \theta_j \end{cases}$$

Notice that *only* if both cards are face up then the value of the cards matters, and the player with the highest card obtains a prize of 10, which is shared equally if the two cards are equal.

In the unique BNE players play D for every but a few of the very highest cards.[8] This may sound surprising at first sight, and this surprise is what drove our choice of the game, in that we aks how a player with limited depth of reasoning — as we defined it — would play this game.

Discussion of the game. This game is suitable to test our proposal of k-level reasoning under private information because the choice between D and U depends on the subjective belief over a single key-parameter, namely α_i.[9] At first sight, action U might seem very tempting because it is a fair split of a prize of 4 rather than 1.[10] In particular, the higher is *your* card, the greater is your expected payoff when actions are $\{U, U\}$. But if you are cautious enough, you know that your rival also understand this, she plays U only with sufficiently high cards, and thus your probability of winning the prize 4 for any given card you draw decreases. The iteration of this reasoning yields to the unique BNE being that D is played with every but possibly the very highest card.

An L0 — according to our definition — has the most simplistic beliefs, which in this game make her believe that $E_i[\alpha_i] = 2$.[11] This belief implies that an L0 plays

[8] A more formal way to see that this is a BNE is as follows. A player's action does not significantly affect her payoff if her rival plays D, and thus we can focus on the right column of the payoff matrix. Additionally, a player's belief over her α increases in her θ, and thus any equilibrium must be of the form "U if my card is above a threshold value" (the threshold might of course be 1 or M too). Thus, say that your rival plays U for all the best $x \in \mathbb{N}_{++}$ cards; that is, for cards $M, M-1, ..., M-x+1$. Then, if your card is $\theta_i = M - x + 1$, by playing U *at best* you tie; that is, your expected payoff of playing U is $\varepsilon + 2/x$, and since a deviation to D gives a payoff of 1, such devation is profitable $\forall x > 2/(1-\varepsilon)$, and thus in a BNE you play D with every but the three very highest cards.

The above is true unless players are too risk-loving, but we neglect this possibility.

[9] According to the players' subjective expectations of α_1 and α_2 (i.e., players' confident of having a high/low chance of winning with their cards), the game boils down to well-known games:

1) a *chicken game* if $\alpha_1, \alpha_2 \in [0, 1)$. In its standard terminology, U corresponds to Swerve, D corresponds to go Straight.

2) an *Active Samaritan's Dilemma* (see Buchanan, 1975) if $\alpha_i \in [0, 1)$ and $\alpha_j > 1$. In its standard terminology, the confident player (in Buchanan's story, the donor) has a dominant strategy U (to help the recipient), and since the non-confident player (the recipient) knows this, she will choose the action D (low effort) which is not what the confident player would like the non-confident player to choose.

3) if $\alpha_1, \alpha_2 > 1$, then the game is a variation of a prisoner's dilemma, where defection by both players ($\{U, U\}$) yields a greater payoff than cooperation by both players ($\{D, D\}$). To the best of our knowledge, this game format has not been named yet, in that it is not interesting *per se*: there is a strictly dominant strategy which maximizes individual and overall payoffs. Thus, virtually every game theoretic prediction would trivially predict $\{U, U\}$.

[10] This is true for risk neutral or not too risk-averse players. We will extensively talk about risk attitude throughout the paper.

[11] A similar but tecnhically more complicated case to analyse is that $\alpha_i = \theta_i - \theta_j$. This case

U regardless of the card she draws; we call this strategy *blind-U*.[12] Action U is arguably an instinctive "first-sight reaction" to the game.[13]

An L1 is one step deeper in her reasoning, in that she computes correctly first-order (and no more) beliefs. Her incapability of correctly computing second-order beliefs makes her regard her rival as an L0, who thus plays *blind-U*. Therefore, an L1 plays U only if her card is any of the top-75% of the cards of the deck (so that $E_i[\alpha_i] \geq 1$), and D otherwise.[14] We call this threshold strategy *Utop75%*.

An L2 is one further step deeper in her reasoning, in that she computes correctly up to the second-order (and no more) beliefs. Equivalently, an L2 best replies to an L1 whose card is uniformly distributed in $\{1, M\}$. Now that the rival plays U only with a top-75% card (rather than *blind-U*), an L1 plays U only if her card is any of the top-75% of the top-75% of the cards of the deck (so that $E_i[\alpha_i] \geq 1$), and D otherwise.[15] Thus, an L2 plays U if her card is any of the top-56%.

Thus, as the depth of reasoning k increases, the above iteration of reasoning, anchored at the L0 level, *monotonically* decreases the threshold above which a player plays U, all the way down to the BNE, which dictates that players play D with any but (possibly) with the very highest card (i.e., $\theta^*_{BNE} = M$).[16] This monotonicity

needs the matrix of payoffs to be normalized by M. We take the simpler path of assuming that α_i is a step function.

[12] The role of ε is as a shortcut to rule out that the $\{D, D\}$ is an equilibrium of the game played by an L0. In fact, we consider $\{D, D\}$ as an implausible equilibrium of the L0 game; several alternative approaches would equivalently rule out $\{D, D\}$ as an L0's equilibrium, such as: (i) the assumption that in case of indifference between U and D a player chooses U, (ii) any of the most common equilibrium refinements, such as trembling hand, strict, payoff-dominance, (iii) it is hard to argue that when we observe action D in reality it is the outcome of indifference between D and U when taking the rival action D as given, and we rather interpret D to be driven by the player's low card, (iv) we believe that a treatment which replace $\alpha_1 = \alpha_2 = 2$ would have virtually all subjects playing U.

Additionally, an L0 with sufficiently high risk-aversion, finds D more profitable. In the lab experiment, this will play a crucial role, but we can easily control for it. In this preliminary discussion, we consider risk-neurality for the sake of the reasoning.

Both these assumptions will be discussed later into details, but dropping them now would confund the main point we want to make.

[13] Action U can alternatively be seen as the "salient action", see Schelling (1960). Finally, a player's limited depth of reasoning could be understood as the outcome of a cost-benefit analysis, where the costs are purely cognitive, see Alaoui and Penta (2016a,b).

[14] Note that this reasoning considers for simplicity the probability of ties to be negligible, or equivalently assumes that M is sufficiently big. In Section 2 we compute the exact threshold, which is $\theta^*_{L1} = \frac{M}{4} + \frac{1}{2}$. That is, for sufficiently high M, an L1 plays U when drawing any of the top-75% of the card.

[15] When your rival plays D, your action does not significantly affect your payoff, and thus this contingency may be neglected. When your card is say the top 74%, against a rival who plays action *Utop75%* you lose almost certainly, and thus deviations to D sufficiently close to the threshold of the L$k-1$ player are profitable for an Lk player.

[16] The parameter choice makes D never a dominant strategy for all θ_i, and thus the threshold of the BNE is interior. Further discussion, analysis and a generalization are in Section 2, but an easy way to see the BNE is that at θ^*_{BNE} the payoff of U and D are equal, i.e. $1 = \frac{\theta^*_{BNE}}{M}2$, so that $\theta^*_{BNE} = M/2$. However, this computation assumes that in case of indifference between D and U a player chooses D, so that a player who draws the very lowest card of the set for which U is better, being indifferent between D and U, chooses D, and thus knows that if she was to deviate to U her expectation of α_i is 0. This assumption is not wlog when θ^*_{BNE} happen to be an integer, which could well be the case, but it simplifies the analysis. See Section 2 for further discussions.

makes the game suitable to identify the depth of reasoning, as defined by our non-equilibrium definition, in a laboratory experiment.[17]

Robustness of L0. In this particular game, plenty of different assumptions on the L0's action or beliefs would yield *blind-U* as the L0's action. For instance, L0 plays *blind-U* for any belief such that $E_i[\alpha_i] \geq 1$, and such belief could be behaviorally motivated in several different ways, since α_i is the oucome of a fair mechanism splitting an overall prize of 4; e.g., taking the simple average of the three possible payoffs values of α_i. Additionally, *blind-U* is the unique strategy under maximin choice, that is, it secures the highest certain prize. Finally, *blind-U* is the best reply against a rival who plays any randomization between D and U which does not depend on the value of the card (as commonly assumed by the standard level-k literature; an L0 uniformly mixes over the whole action space). Thus, upon controlling for risk-aversion, we are confident that *blind-U* is a suitable behavior for the instinctive L0 players. Notice that among the goals of the paper there is *not* to propose a better L0 or to test which L0 fits best the data. Our L0 comes naturally with our definition of k as the depth of reasoning concerning the *strategic role of information* in the game.

Best replying. Admittedly, best responding in our game is less straightforward than in other games which have been used to test level-k (such as the 11-20 money request game of Arad and Rubinstein, 2012), and the reason is inherent of private information game, which, as opposed to complete information games, make players bear the extra cognitive burden of understanding the role of information in affecting actions and beliefs. This is why we obviously do not expect subjects to act according to a precise threshold strategy as our non-equilibrium computations predict. We rather built the game so as to have the properties that the threshold monotonically decreases with the player's depth of reasoning. Thus, it is perhaps more cautious to interpret our "k" in an ordinal, rather than cardinal, fashion. In fact, comparing our k with the standard complete information k would be a structural mistake, since they measure things for which the cognitive burden to push one's own reasoning to a certain level k might greatly change.

Risk. Behaviors naturally depend on the players' attitude towards risk. However, we do not regard this as an issue to solve for at least three reasons. First, it is easy to elicit risk and control for it in the laboratory, for example making subjects choose between 1€ and an even lottery between 0€ and 4€. Second, at any depth of reasoning k, the same amount of risk is involved, because — at any depth — the subject compares 1€ and the 0-2-4€ lottery where the weights are exclusively determined by the expected ranking between my card and my rival's one; in other words, for any threshold-strategy of my rival (and thus for any k), I always best reply by trimming "from the bottom" the same share of the cards from which I play U (this share is 25% in our numerical), and this share depends on my risk propensity,

If one were to assume that in case of indifference between D and U a player chooses U, then $1 + \frac{M-\theta^*_{BNE}}{M}\frac{1}{M-\theta^*_{BNE}}2 = \frac{\theta^*_{BNE}}{M}2$, so that $\theta^*_{BNE} = 1 + M/2$.

[17] Likewise, in a p-beauty contest, the chosen number monotonically decreases with the player's depth of reasoning.

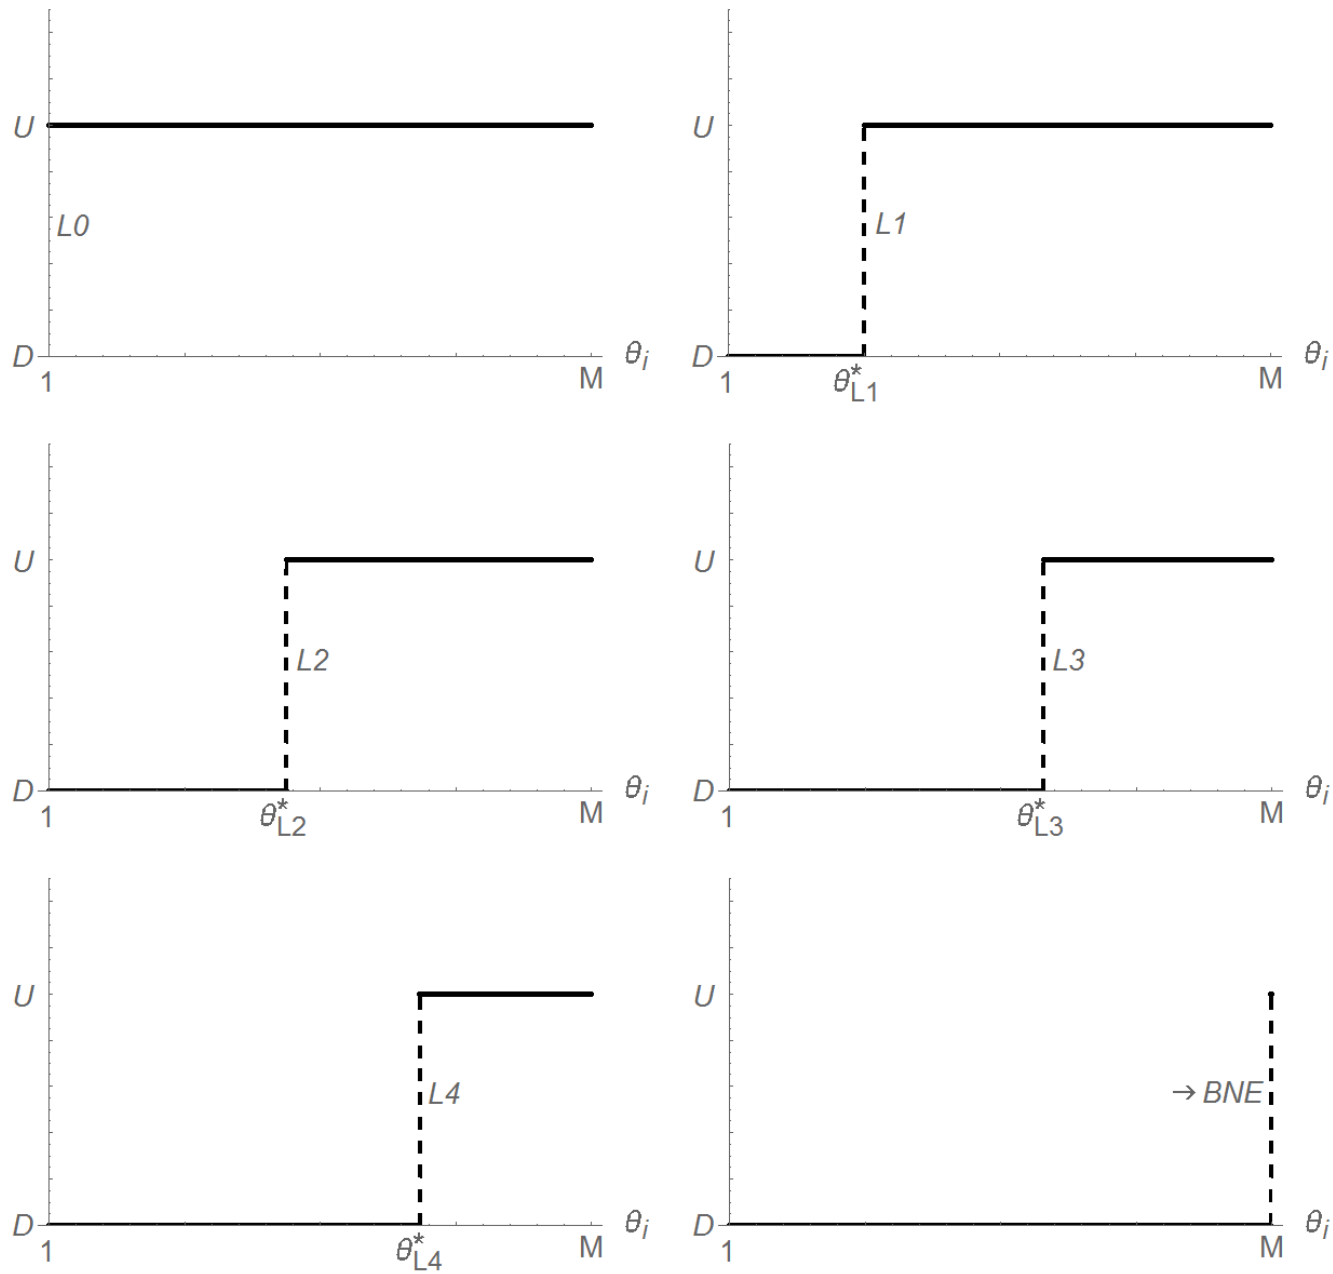

Figure 1: Evolution of Lk strategy in the Up-or-Down game from L0 to the BNE.

but *not* on k. Third, payoff uncertainty is an inherent feature of game of private information on types (i.e., some payoff-relevant parameters) where, on the contrary of complete information games, the action depends on beliefs over types which are inherently distributions over the type-space.

1.2 L0's behavior as a cognitive bias

> *"perch' elli 'ncontra che più volte piega // l'oppinïon corrente in falsa parte, // e poi l'affetto l'intelletto lega. " Dante Alighieri, Divina Commedia, Paradiso 13, 118-121.*
>
> *English translation: "opinion—hasty—often can incline to the wrong side, and then affection for one's own opinion binds, confines the mind."*

Our assumption is that L0-players ignore the possibility of heterogeneity of types and believe that their rival is simply of her same type. The discussion of such assumption changes according to the interpretation that one gives to players' types. All the points made below in some sense rely on people's cognitive misperceptions or heuristics. The most direct available heuristic - measured as how readily a particular idea comes to mind - is one own observation, which itself motivates our choice of L0-players' behavior.

According to the **representativeness heuristic**, the estimates of subjects are insensitive to prior probabilities when subjects are given some piece of information, as shown originally by Kahneman and Tversky (1973). In their experiment, subjects are given a description of a person named Jack[18], and they are told that Jack was randomly chosen from population with a given proportion of lawyers and engineers. In one treatment, the proportion of engineers was 0.3, while in the other treatment the proportion of engineer was 0.7. When asked to assess the probability that Jack is an engineer, subjects of the two treatments estimated essentially the same probability. They relied on whether the description they read is representative of the stereotype of the engineer, rather than on the prior probability. This is what Kahneman and Tversky call the representativeness heuristic. Completely neglecting the prior probability is, to the extreme, what the L0-players do in our proposal of non-equilibrium concept.[19]

The **hindsight bias** is the subjects' inability to accurately remember their prior expectations after observing some new piece of information. In particular, subjects tend to ex-post overestimate the ex-ante probability of the observed piece of information. This phenomenon was initially found by Fischhoff (1975), and the subsequent literature proved its robustness.[20] For instance, in Loewenstein, Moore

[18]Namely, "Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing, and mathematical puzzles."

[19]For numerous other examples and applications of the representativeness heuristic see Kahneman et al. (1982), chapters 2-6.

[20]See Blank et al. (2007) for an excellent survey of the hindsight bias research. A theoretical model on the effect of hindsight bias voters on political elections is Schuett and Wagner (2011). Biais and Weber (2009) characterize the effect of hindsight bias in the financial market on investment and trading.

and Weber (2006) subjects are shown two pictures and they have to spot the one difference. Some subjects are informed and some others are not about the difference. When informed subjects are incentivized to guess the fraction of uninformed subjects who would be able to spot the difference, they guessed 58%, while the true fraction was 20%. Another example is that before George W. Bush had initiated military action in response to the Iraqi invasion of Kuwait, opinions about the rightness of military action was about equally distributed between in favor and against, but after observing the success of such mission, 4 to 1 recalled that they were in favor of military action even before observing its outcome (see Mueller, 1994). Such subjects' overestimation with hindsight of the ex-ante probability of the observed piece of information can be stretched in its extreme to perceiving the prior as a degenerate distribution giving all the weight to one's own observation, and thus to the L0-players' behavior.

The fact of placing too much weight on one's own type can be the outcome of what psycologists call the **confirmation bias**; that is, the tendency to notice and actively look for what confirms one's preexisting ideas, and to neglect and not look for alternative possibilities. For instance, say that a person's type is whether she believes or not in a correlation between the department where a journal is edited and the editors' leniency toward papers submitted by authors from that same deparment. Then, according to the confirmation bias, subjects tend to place greater importance to evidences that upholds their type and to ignore evidences that challenge their type; subjects "notice more" when a paper is published by authors affiliated at the journal's editorial headquarter than when a paper's authors have a different affiliation. As a result, people tend to overestimate the number of people who share their same type.[21]

Also closely related to the hindsight bias is the **information projection**. While the hindsight bias is the overweighting of one's own information in the prior, the information projection is the direct projection of one's own information onto others. As reported by Madarasz (2012), "people too often think as if others knew what they did, and too often act as if others could guess their private information correctly". In his paper, Madarasz proposes a model where players systematically overestimate the access of other players to their own private information. This idea has the flavor of the behavior of L0-players, who entirely project their private information to the other players.

When "type" is interpreted as norms, preferences, habits, or opinions, then overestimating the degree to which one's own information is shared by other people takes the name of **false consensus effect**. Because *you* think that A is better than B, you assume that the vast majority of other people also feel the same way you do. This cognitive bias tends to lead to the perception of a consensus that does not exist, a "false consensus". An american going to Australia for the first time will tend to incorrectly believe that "up" of a toggle switch turns on the light, rather than off. The experiment carried out in Ross et al. (1977) reports that college students who preferred brown bread over white bread estimated that 52.5% of all college students preferred brown bread, whereas college students who preferred white bread estimated that only 37.4% of all college students preferred brown bread.

[21] For an excellent review of confirmation biases, see Nickerson (1998).

This suggests that some people who observe that brown bread is better for themselves, tend to believe that others also think the same way that they do. A more economical example of false consensus effect is as follows. Say that a person's type is her elasticity of relocating to another country to the tax burden, and say that the tax burden increases by 3%. A person with sufficiently low elasticity will not relocate to another country as a result of such a tax increase, and will also tend to erroneously believe that such a tax increase does not have an effect on emigration because it does not have an effect on her own emigration choice.[22]

The above definitions span antropology, sociology and psycology, and are highly intertwined, and often overlap. Other branches of economics have recently highlighted the importance of agents' limited capacity of processing information. For instance, a sound branch of dynamic programming started with the seminal contribution of Sims (2003), who shows that accounting for limited information-processing capacity of agents can explain several observed macroeconomic behaviors. One way to intrepret the instinctive behavior of our L0-players is that they have limited capacity to process the information contained in the prior distribution of types.

2 The generalized Up-or-Down-game

In this section we generalize the payoff matrix of the Up-or-Down game so as to build a better understanding and hence to choose carefully the payoffs in the lab experiment treatments. Two players privately observe their types θ_i with $i = 1, 2$, which is drawn from a commonly known $\theta_i \sim U\{1, M\}$ with M "big".[23] I call u (i.e., ε in the Introduction) the certain payoff of action U, and I call d the payoff a player obtains if her action is D and her opponent's one is U. Thus, the payoff matrix is

	D	U
D	$(0,0)$	(d,u)
U	(u,d)	$(u+\alpha_1, u+\alpha_2)$

where,

$$\alpha_i = \begin{cases} 2\pi & \text{if } \theta_i > \theta_j \\ \pi & \text{if } \theta_i = \theta_j \\ 0 & \text{if } \theta_i < \theta_j \end{cases}$$

In what follows we will make three assumptions *(A1)-(A3)* on the quadruple (u, d, π, M), which are clearly satisfied by the numerical example in the Introduction.

The first assumption we make is

$$\textit{(A1)}\text{: } 0 < u < d$$

otherwise, regardless of π, action U would be strictly dominant regardless of beliefs and depth of reasoning.

[22] See Mullen et al. (1985) for a meta-analysis of the false consensus effect.

[23] The reader would have to wait till "**Remark on** M" below for the meaning of "big".

Symmetrically, in order to avoid D to be too tempting (dominant) we assume that

$$(A2)\text{: } d < \min\{u+\pi, 2\pi\}$$

Condition $d < u + \pi$ is necessary to guarantee that the complete information game that an L0 plays (i.e., $\alpha_1 = \alpha_2 = \pi$) has the unique equilibrium $\{U, U\}$, instead of two equilibria; namely, $\{U, D\}$ and $\{D, U\}$. Thus, under *(A2)*, U is tempting enough so as to unambiguously make an L0 play *blind-U*.

Condition $d < 2\pi$ guarantees that D is not the dominant action $\forall k$ when u is arbitrarily small (like in the example of the Introduction).

Analysis of L1. An L1 of a certain type θ_i best replies to a rival whose: i) type is $\theta_j \sim U\{1, M\}$, and ii) action is U. Then, the L1 plays U rather than D iff

$$\begin{aligned} u + E_i[\alpha_i \quad & | \quad \theta_i, k_i = 1] \geq d \\ \iff \quad & u + \pi \underbrace{\frac{1}{M}}_{\Pr\{\theta_i=\theta_j\}} + 2\pi \underbrace{\left(\frac{\theta_i - 1}{M}\right)}_{\Pr\{\theta_i>\theta_j\}} \geq d \\ \iff \quad & \theta_i \geq \frac{M(d-u)}{2\pi} + \frac{1}{2} \end{aligned} \tag{1}$$

And if we define $\theta^*_{L1} = \left\lceil \frac{M(d-u)}{2\pi} + \frac{1}{2} \right\rceil$, where $\lceil x \rceil$ is the smallest integer greater than or equal to x, then an L1 plays[24]

$$\begin{cases} U \text{ if } \theta_i \geq \theta^*_{L1} \\ D \text{ if } \theta_i \leq \theta^*_{L1} - 1 \end{cases} \tag{2}$$

Notice that in the argument given in the Introduction, we neglected for simplicity the probability of ties, which yield term "$+\frac{1}{2}$" in (1). Thus, if we did not neglect this probability, the threshold card for an L1 would be shifted by one card as opposed to *Utop75%* if term "$+\frac{1}{2}$" does not affect the L1's strategy (i.e., mapping from θ_i to $\{U, D\}$), or equivalently if $M \neq 4n + 2$ with $n \in \mathbb{N}_+$.[25] In other words, if $M \in \{6, 10, 14, 18, ...\}$ then the true θ^*_{L1} (i.e., the one which does not neglect the probability of ties) shifts the threshold by one card as opposed to strategy *Utop75%*.

Analysis of Lk. Since from L2 onwards the player best responds to a threshold strategy (mixture of D and U), we can already provide the general analysis of Lk. An Lk of type θ_i best replies to a rival whose: i) type is $\theta_j \sim U\{1, M\}$, and ii)

[24] Leaving the threshold as a possibly non-integer number would also yield to an unambiguously defined strategy, but we round the threshold in order to be able to use θ^*_{L1} to compute the exact probability that an $L1$ plays D when computing the L2's action.

[25] The fact that $M = 4n + 2$ with $n \in \mathbb{N}_+$ shifts the threshold can be seen in the table below.

M	$\frac{1-\varepsilon}{4}M + \frac{1}{2}$	$\theta^*_{L1}(Utop75\%)$
3	1.249	1.5
4	1.499	1.75
5	1.749	2
6	1.999	2.25
7	2.249	2.5
8	2.499	2.75

action is U iff $\theta_j \geq \theta^*_{Lk-1}$. An $Lk-1$ plays, as in (2), plays U with probability $(M-\theta^*_{Lk-1}+1)/M$. Then the Lk plays U rather than D iff

$$
\begin{aligned}
u + \frac{M-\theta^*_{Lk-1}+1}{M} E_i[\alpha_i \quad &| \quad \theta_i, k_i = k] \geq \frac{M-\theta^*_{Lk-1}}{M} d \\
&\iff \quad u + \frac{M-\theta^*_{Lk-1}+1}{M} \left[\underbrace{\frac{1}{M-\theta^*_{Lk-1}+1}}_{\Pr\{\theta_i=\theta_j\}} \pi + \underbrace{\frac{\theta_i-\theta^*_{Lk-1}}{M-\theta^*_{Lk-1}+1}}_{\Pr\{\theta_i>\theta_j\}} 2\pi \right] \geq \frac{M-\theta^*_{L}}{M} \\
&\iff \quad u + \frac{\pi}{M} + \frac{2\pi}{M}\left(\theta_i - \theta^*_{Lk-1}\right) \geq \frac{M-\theta^*_{Lk-1}}{M} d \\
&\iff \quad \theta_i - \theta^*_{Lk-1} \geq \frac{M(d-u) - d\theta^*_{Lk-1} - \pi}{2\pi} \\
&\iff \quad \theta_i \geq \theta^*_{Lk-1} + \frac{M(d-u) - d\theta^*_{Lk-1} - \pi}{2\pi}
\end{aligned}
$$

And if we define $\theta^*_{Lk} = \left\lceil \theta^*_{Lk-1} + \frac{M(d-u)-d\theta^*_{Lk-1}-\pi}{2\pi} \right\rceil$, then an Lk plays

$$
\begin{cases} U \text{ if } \theta_i \geq \theta^*_{Lk} \\ D \text{ if } \theta_i \leq \theta^*_{Lk} - 1 \end{cases}
$$

First, the fixed point at $k \to \infty$ requires the right-hand side of (3) to equal 0, or equivalently,

$$
\theta^*_{L\infty} = \frac{d-u}{d} M - \frac{\pi}{d} \tag{4}
$$

As $M \to \infty$, the threshold converges to playing U only for the top $\left(\frac{d-u}{d}\right)\%$ cards.[26] Thus, in the numerical example of the Introduction where u is arbitrarily small, a player playing according to the BNE plays U only if the card is the very highest one of the deck.

Second, the process determining the thresholds is a linear first-order difference equation with constant coefficient of the form $x_t = ax_{t-1} + b$ with $a = \frac{2\pi-d}{2\pi}$ and $b = \frac{M(d-u)-\pi}{2\pi}$, see (3). Under *(A2)*, $a \in (0,1)$ and thus the process convergece monotonically.To further guarantee a well-behaved problem, we assume what is necessary for $\theta^*_{L\infty} > \theta^*_{L1}$, namely,

$$
\begin{aligned}
\frac{d-u}{d} M - \frac{\pi}{d} \quad &> \quad \frac{M(d-u)}{2\pi} + \frac{1}{2} \\
&\iff \quad (2\pi - d)(d-u)M > (d+2\pi)\pi \\
&\iff \quad \textit{(A3)}\text{:}\ \frac{2\pi-d}{2\pi+d}(d-u)M > \pi
\end{aligned}
$$

So that *(A3)* guarantees $\theta^*_{L\infty} > \theta^*_{L1}$. Notice that by *(A2)*, $\frac{2\pi-d}{2\pi+d} \in (0,1)$, and thus *(A3)* implies $(d-u)M > \pi$, which guarantees $\theta^*_{L1} > 1$ — see (1). Since $\theta^*_{L\infty} < M$ —

26Notice that $\frac{d-u}{d} \in (0,1)$ by *(A1)*.

see (4) — assumptions *(A2)-(A3)* guarantee $1 < \theta^*_{L1} < ... < \theta^*_{Lk} < ... < \theta^*_{L\infty} < \dot{M}$, and thus the game is "well-behaved". Notice that under *(A1)-(A2)*, there is always an M big enough so that *(A3)* holds.

Remark on M. In the Introduction we used $u = \varepsilon > 0$ arbitrarily small, $d = 1$ and $\pi = 2$, but we also assumed that the deck is sufficiently "big" so as to neglect in our reasoning for simplicity the probability of ties, which are rather crucial for small decks. In fact, if $M = 1$ then $\alpha_i = \alpha_j = \pi\ \forall k$ and thus k does not affect the strategies. Additionally, while *(A1)* and *(A2)* do not depend on M, *(A3)* requires the deck to be composed of at least 4 cards (in particular, $M > 3.\bar{3}$). In particular, for any triple (u, d, π) satisfying *(A1)* and *(A2)* there is always an M big enough so that *(A3)* hold. Finally, notice that the analysis for a continuous distribution of types is equivalent to the limiting case $M \to \infty$. However, the fact that M has to be big enough is only needed in the generalized version of the game presented here, and not in the Introduction, where $M \geq 3$ suffices for *(A3)*.

Notice that the entire analysis can be easily replicated for any distribution of types (non-uniform) over a compact space, where all the thresholds of the strategies would be specified in quantiles.

3 Laboratory Experiment

Consider the generalized payoff matrix of the Up-or-Down game.

	D	U
D	$(0, 0)$	(d, u)
U	(u, d)	$(u + \alpha_1, u + \alpha_2)$

If we normalize $\pi = 2$ as in the Introduction, the degrees of freedom are on the choice of two key-parameters, namely u and d. The only constraint in their choice is to respect *(A1)-(A3)*, where *(A3)* simply says that the number of cards is sufficient, while *(A1)-(A2)* impose that: $u < d < \min\{u + 2, 4\}$. We propose six treatments which are informative on different players' behaviors:

T1: $u = 0, d = 1$
T2: $u = 0, d = 0.3$
T3: $u = 0, d = 1.7$
T4: $u = \varepsilon > 0, d = 1$
T5: $u = 0.5, d = 1$
T6: $u = 2, d = 3$

Several tradeoffs emerge in the choice of u and d, and these tradeoffs motivate our choice of treatments. In particular:

1. $u = 0$ makes best-responding very natural. In fact, think about best replying to $Utop75\%$ in T1. If the player's card is at (or arbitrarily close to) the top 75%, then $E[\alpha] \approx 0$, and thus U gives $u = 0$, while D is 3 times more likely to give $d = 1$ than 0, hence it gives in expectation $0.\bar{6}$, and thus the

threshold increases as k increases. Thus, T1-3 are our favorite treatments for how natural it is to best reply. However, setting $u = 0$ has potentially two drawbacks: i) a reasoning such as "joy of winning some positive prize" might kick-in, and ii) indifference between D and U formally yield an implausible L0 equilibrium $\{D, D\}$, which we already discussed in the Introduction. A comparison between T1 and T4 will shed light on such implausibility.

2. however, the higher is u, the more natural is U to be L0. In fact, playing U is a tempting instinctive reaction to the game for those players who do not want to or are incapable of engaging in strategic thinking, in that it secures the highest certain prize u. Thus, a comparison of how many subjects play *blind-U* between T1 and T5 would shed light on this aspect.

3. increasing both u and d by the same amount anchors U as an L0 strategy for higher risk-aversion levels. Thus, T6 is our favorite treatment to avoid risk-aversion concerns, and its comparison with T1.

4. the lower is u, the more the threshold in the strategies of level k's spans all the way down to playing always D for players with high depth of reasoning, (the BNE is $\theta^*_{BNE} = \frac{d-u}{d}M$ for sufficiently high M), and thus low u's provide more variation of behavior and better identification of k.

We take an agnostic approach and run six treatments, in that this game has not been tested before in the laboratory, and thus we do not know the comprehension level of subjects of the game in general.[27]

The experiment is programmed using z-Tree (Fischbacher 2007) and run at the Max Planck Institute laboratory in Munich, Germany. Subjects will be recruited from the student body of Munich universities using ORSEE (Greiner 2004). We will admit 24 subjects to each session. Each subject participated in exactly one of the treatments outlined above. Subjects will play the game described in Section 2 repeatedly (20 independent rounds in total), but in each repetition subjects will be randomly rematched. The subjects will not be given specific information about the precise nature of matching mechanism other than that they would be randomly rematched between rounds and that they will never play against the same opponent more than once, so as to obtain a larger number of independent observations where learning about rivals plays no role.

At the beginning of each session, we distribute and read out loud the instructions (see the Appendix), and subjects have to complete a quiz to make sure they understood the Up-or-Down game. The numer of rounds is set to 20 so as to have a sufficient number of observations to identify with confidence the threshold below (above) which the subject play U (D). At the end of each round, subjects are disclosed their rival's action and their own payoff. Subjects are not told their rival's card, otherwise this might trigger updating of their beliefs about their rival's depth of reasoning, and this might lead them to infer the average depth of reasoning of

[27] Another possible treatment is to compare small M with large M: increasing M makes the probability of ties go to zero, and thus the L0 reasoning might be less likely to naturally kick-in. Technically, $M \to \infty$ is equivalent to a continuous type and action space.

others, and thus they might react accordingly in subsequent periods. We do not want this to affect subjects' behavior. Additionally, at the end of each round (and thus also at the end of the session) an explanation sheet is given to invite subjects to explain the reasoning they followed. Subjects are paid one (or more) randomly selected round. They are told which round has been randomly selected only at the end of the experiment. At the end of the session, we elicit subject's risk-aversion by making them choose between lotteries and certain equivalents. In particular, the certain equivalent is u (crucial in T6) and the lottery is between 0 and d. Subjects will also be asked what they believe on others' choice in this lottery, so that their choice is not affected by beliefs of other's risk aversion).

A strategy for a player is a mapping from $\{1, M\}$ to a binary action $\{D, U\}$. In each round we observe only one point of this strategy. Thus, we need several rounds in order to have a better identification of the strategy played by the players. Observing D for a given card would give us a *lower-bound* in the depth of reasoning of the player, whereas observing U for a given card would give us a *upper-bound*. The smallest action space in order to guarantee perfect identification of k for every possible type is when the behavior changes for every pair $\{k, M\}$ with $k \to \infty$, and this means to have a strategy space of dimension $|M| \times \infty$, rather than the simple binary space we used. Thus, there is a trade-off between simplicity of the game and identification power.

4 Application of the non-equilibrium concept to simple Bayesian games

The game proposed in this paper is intended only to provide a clear-cut, albeit rather special, framework which is suitable to study our non-equilibrium proposal. It is also of interest to see what are the predictions in "standard" games, which is what we do in this Section. We find that several "anomalies" might happen. In Cournot games, the infinite iteration of reciprocal belief formation ($k \to \infty$) does converge to the BNE, but non-monotonically. In Contests, there is convergence, which is monotone from $k = 1$ onwards. In the Market entry game, the convergence to the BNE is lost. In the sealed-bid irst-price auction, convergence to the BNE already occurs at L1. This highlight an important discontinuity. In particular, this shows that, in general, the BNE is not a suitable tool to analyze equilibrium behaviors under finite depth of reasoning.[28]

4.1 Cournot

Two firms indexed by $i \in \{1, 2\}$ compete a lá Cournot with linear demand and linear cost of quantity,

$$u_i(q_1, q_2) = q_i(2 - q_1 - q_2) - \theta_i q_i$$

[28] One of the first paper pointing out a discontinuity of the infinite iteration of reciprocal expectation was Rubinstein (1989).

where the marginal cost θ_i is an iid draw from a $U[0,1]$. Firm i privately observes their θ_i but does not observe θ_{-i}.

An L0-firm of certain type θ_i believes that her opponent is also of type θ_i with identical beliefs (i.e., in complete information). Hence, an L0-firm of certain type θ_i chooses $q_{\theta_i}^{L0}$ which solves the following FOC[29]

$$2 - 2q_{\theta_i}^{L0} - q_{\theta_i}^{L0} = \theta_i$$

and thus

$$q_{\theta_i}^{L0} = \frac{2-\theta_i}{3}$$

An L1-firm of certain type θ_i *correctly* believes that her opponent's type is uniformly distributed in $[0,1]$, but also believes that her opponent is L0. Thus, an L1-firm of certain type θ_i chooses $q_{\theta_i}^{L1}$ which solves the following FOC

$$2 - 2q_{\theta_i}^{L1} - \int_0^1 \frac{2-x}{3} dx = \theta_i$$

and thus

$$q_{\theta_i}^{L1} = \frac{3-2\theta_i}{2}$$

Subsequent iterations of this procedure yield

$$\begin{aligned} q_{\theta_i}^{L2} &= \frac{1-\theta_i}{2} \\ q_{\theta_i}^{L3} &= \frac{14-4\theta_i}{8} \\ &\ldots \end{aligned}$$

More in general,

$$2 - 2q_{\theta_i}^{Lk} - \int_0^1 q_{\theta_i}^{Lk-1} dx = \theta_i$$

where, taking the linear form $q_{\theta_i}^{Lk} = ax+b$, the fixed point is when $q_{\theta_i}^{Lk} = q_{\theta_i}^{Lk-1} = ax+b$, and thus

$$\begin{aligned} 2 - 2(a\theta_i + b) - \int_0^1 (ax+b)dx &= \theta_i \\ &\iff 2 - 2a\theta_i - 2b - \left(\frac{a}{2} + b\right) = \theta_i \\ &\iff 2 - 3b - \frac{a}{2} = \theta_i(1+2a) \end{aligned}$$

[29] Throughout the paper we assume type-symmetry. That is, identical types choose identical strategies.

Since it has to hold $\forall\theta_i$, then $a = -\frac{1}{2}$, and thus $b = \frac{3}{4}$. It is easy to calculate the BNE of this game is indeed $q^{BNE} = \frac{3-2\theta_i}{4}$.[30] Thus, in such Cournot game there is convergence, but non-monotone, to the BNE as $k \to \infty$.

4.2 Contest

Two contestants indexed by $i \in \{1, 2\}$ play a Tullock contest with linear cost of effort and a prize of value 1,

$$u_i(e_1, e_2) = \frac{e_i}{e_i + e_{-i}} - \theta_i e_i$$

where θ_i could be $\theta_L = 2$ or $\theta_H = 1$ with equal probabilities. Contestant i privately observes her θ_i but does not observe θ_{-i}.

An L0-player of certain type θ_i believes that her opponent is also of type θ_i with identical beliefs (i.e., in complete information). Hence, an L0-player of type θ_i chooses the following $e^{L0}_{\theta_i}$

$$e^{L0}_{\theta_i} = \frac{1}{4\theta_i} \quad \Rightarrow \begin{cases} e^{L0}_{\theta_L} = 0.125 \\ e^{L0}_{\theta_H} = 0.25 \end{cases} \tag{5}$$

An L1-player of certain type θ_i *correctly* believes that her opponent's type is θ_L or θ_H with equal probabilities, but also believes that her opponent is L0. Thus, an L1-player of certain type θ_i chooses $e^{L1}_{\theta_i}$ solving the following FOC

$$\frac{1}{2}\frac{e^{L0}_{\theta_L}}{(e^{L1}_{\theta_i} + e^{L0}_{\theta_L})^2} + \frac{1}{2}\frac{e^{L0}_{\theta_H}}{(e^{L1}_{\theta_i} + e^{L0}_{\theta_H})^2} = \theta_i \tag{6}$$

where $e^{L0}_{\theta_L}$ and $e^{L0}_{\theta_H}$ are defined in (5). One could solve (6) and find that

$$\begin{cases} e^{L1}_{\theta_L} \approx 0.11666 \\ e^{L1}_{\theta_H} \approx 0.237834 \end{cases}$$

An L2-player of type θ_i *correctly* believes that her opponent's type is an θ_L or an θ_H type with equal probabilities, but also believes that her opponent is L1. Thus, one could compute the equilibrium and find

$$\begin{cases} e^{L2}_{\theta_L} \approx 0.11789 \\ e^{L2}_{\theta_H} \approx 0.235727 \end{cases}$$

The BNE is

$$\begin{cases} e^{BNE}_{\theta_L} = \frac{17}{128} = 0.1328125 \\ e^{BNE}_{\theta_H} = \frac{17}{64} = 0.265625 \end{cases}$$

And one could verify that there is convergence to the BNE. However, this convergence, except from L0 to L1, is monotone in this class of games.

[30] If the distribution of types was binary, with $\theta_H = 0$ with probability p and $\theta_L = 1$ with probability $1 - p$, perfect convergence to the BNE occurs already at L1.

In fact, under such distribution of types, one would obtain

$$\begin{aligned} q^{L0}_{\theta_H} &= \frac{2}{3} \text{ and } q^{L0}_{\theta_L} = \frac{1}{3} \\ q^{L1}_{\theta_H} &= \frac{5-p}{3} \in \left[\frac{2}{3}, \frac{5}{6}\right] \text{ and } q^{L1}_{\theta_L} = \frac{2-p}{3} \in \left[\frac{1}{6}, \frac{1}{3}\right] \end{aligned}$$

and these last quantities correspond to the BNE.

4.3 All-pay auction

Two contestants indexed by $i \in \{1, 2\}$ play a standard all-pay auction with linear cost of effort, private values, and valuations v_i are independently and uniformly distributed on $[0, 1]$.

BNE. It is well known that the unique symmetric BNE is the bidding function $b(v_i) = \frac{v_i^2}{2}$ (see e.g. Wolfstetter, 1999, Chapter 8.2).[31]

L0-player. An L0-player with type v_i plays as in a complete-information symmetric-type all-pay auction. It is well known that the only equilibrium is in mixed-strategies, where the L0-player bids uniformly in $[0, v_i]$ (see Baye, Kovenock and de Vries, 1996).

Remark: at first sight, one might think that the L0 is more difficult than the BNE, but this is true only ex-post. That is, the difficulty to compute an equilibrium and how difficult the equilibrium turns out to look like are very different concepts, and our k captures the former.

L1-player. Let us see what is the BR of an L1-player, called player i, to an L0-player, called player j. Suppose that player L1 bids $b_i(v_i)$ where v_i is his valuation. Then his probability of victory is the probability that $b_i(v_i)$ is greater than a $U[0, v_j]$ where $v_j \sim U[0, 1]$. Thus, his probability of victory equals

$$b_i(v_i) + [1 - b_i(v_i)] \int_{b_i(v_i)}^{1} \frac{b_i(v_i)}{v_j} dv_j$$

The first term $b_i(v_i)$ accounts for the cases when the extraction of $v_j \sim U[0, 1]$ is *below* the bid $b_i(v_i)$ so that i wins for sure, and the second term is the probability of victory of i when the extraction of $v_j \sim U[0, 1]$ is *above* the bid $b_i(v_i)$ so that i does not win for sure. The integral is the probability that player i wins for all possible $v_j > b_i(v_i)$. Thus, the expected utility of the L1-player i who bids $b_i(v_i)$ is,

$$v_i \{b_i(v_i) - [1 - b_i(v_i)] \, b_i(v_i) Log[b_i(v_i)]\} - b_i(v_i)$$

Mathematica shows its quasi-concavity, thus we focus here on the FOCs with respect to $b_i(v_i)$, shortly called b below,

$$\begin{aligned} v_i \left\{ b' - [1 - 2b] \, Log[b] b' - \frac{b - b^2}{b} b' \right\} &= b' \\ b - (1 - 2b) Log[b] &= 1/v_i \end{aligned} \tag{7}$$

which implies that $b_i(v_i) = 0$ when $v_i = 0$, and $b_i(v_i) \to \bar{b}_{L1} \cong 0.235611$ when $v_i \to 1$. Also, $\bar{b}_{L1}$ is the maximum bid of an L1-player.

L2-player. Computing the strategy of an L2-player is significantly harder, because it is the BR to an L1-player, whose bid is implicitely defined by (7) where $v_i \sim U[0, 1]$. If she bids $b_i(v_i)$, any bid greater than $\bar{b}_{L1}$ is strictly dominated, because it does not increase the probability of winning, but it costs more than $\bar{b}_{L1}$.

[31] https://www2.wiwi.hu-berlin.de/institute/wt1/staff/wolfstetter/chap-08-new.pdf formula (8.26)

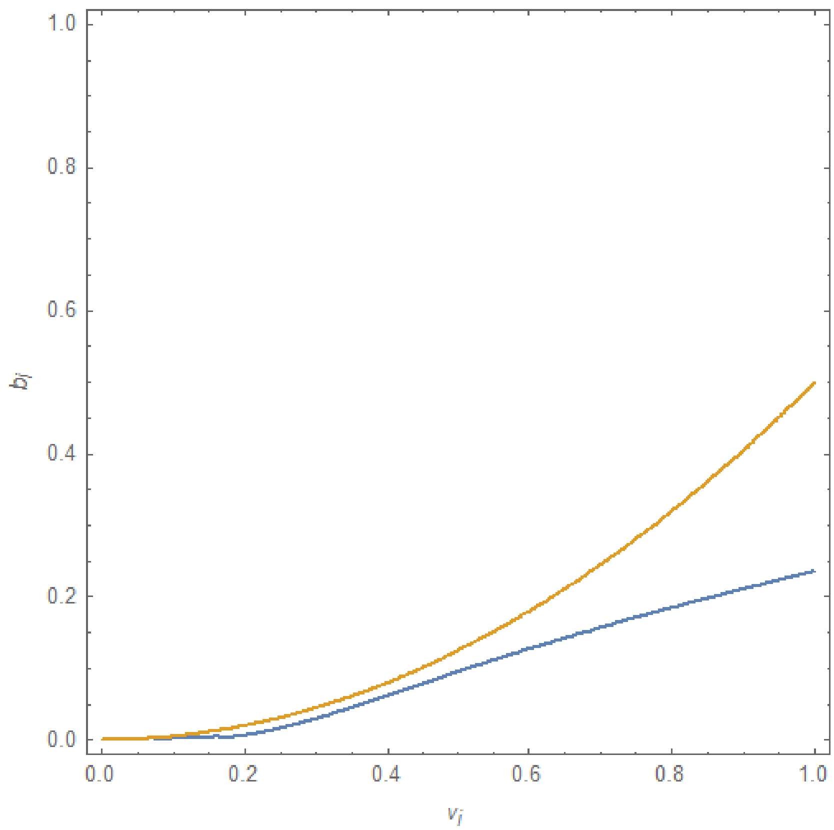

Figure 2: The orange line is the BNE and the blue line is the L1-player's bidding function.

4.4 Market entry

Consider a two-player market entry game, where a player who does not enter obtains a payoff of 0, when both players enter each obtains a payoff of -1, and when only player i enters, she obtains a payoff of θ_i, which is her private information. Thus, the payoff matrix is as follows.

	Enter	Don't enter
Enter	(-1,-1)	(θ_1,0)
Don't enter	(0,θ_2)	(0,0)

Types θ_i are distributed as follows.

$$\theta_i = \begin{cases} \theta_H & \text{w. prob. } \frac{1}{2} \\ \theta_L & \text{w. prob. } \frac{1}{2} \end{cases}$$

where $\theta_H > \theta_L > 1$.

An L0-player of certain type θ_i believes that her opponent is also of type θ_i with identical beliefs (i.e., in complete information). Hence, an L0-player of type θ_i enters with probability $\frac{\theta_i}{\theta_i+1}$.[32]

An L1-player of certain type θ_i *correctly* believes that her opponent's type is θ_L or θ_H with equal probabilities, but also believes that her opponent is L0. Thus, an L1-player expects her opponent to enter with probability $\bar{p} = \frac{1}{2}\frac{\theta_H}{\theta_H+1} + \frac{1}{2}\frac{\theta_L}{\theta_L+1}$. If the L1-player is of type θ_i, she enters if $\bar{p}(-1) + (1-\bar{p})\theta_i > 0$, which holds (does not hold) when $\theta_i = \theta_H$ ($\theta_i = \theta_L$). Thus, high-L1-players enter and low-L1-players do not enter.

An L2-player of type θ_i *correctly* believes that her opponent's type is an θ_L or an θ_H type with equal probabilities, but also believes that her opponent is L1. Thus, an L2-player expects her opponent to enter with probability $\frac{1}{2}$, and since $\frac{1}{2}(-1) + \frac{1}{2}\theta_i > 0 \ \forall \theta_i \in \{\theta_H, \theta_L\}$, an L2-player enters regardless of her type. From Lk with $k \geq 3$ players do not enter if k is odd and enter if k is even. Therefore, Lk's action alternate between the two pure-strategy equilibria of the original game. Note that the original game has also a BNE where types θ_H enter with certainty and types θ_L enter with probability $\frac{\theta_L-1}{\theta_L+1}$.

4.5 First price auction

Two bidders indexed by $i \in \{1, 2\}$ simultaneously bid in a standard first-price sealed-bid auction (FPA) for an item whose value is θ_i to bidder i. Values θ_i's are iid draw from a distribution $U[0, 1]$. Bidder i privately observes her θ_i but does not observe θ_{-i}.

It is well-known that the BNE is to bid $\theta_i/2$, that is, half of the own privately observed valuation.[33]

[32] In fact, the indifference condition reads

$$p(-1) + (1-p)\theta_i = 0$$

[33] E.g., see Wolfstetter E (1999), Chapter 8, or even easier: https://en.wikipedia.org/wiki/First-price_sealed-bid_auction

An L0-bidder of certain type θ_i believes that her opponent is also of type θ_i with identical beliefs (i.e., in complete information). This makes an L0-bidder behaves as in a common-value full-information FPAs, it is well-known that in equilibrium bidders bid their own valuation.[34] Thus, all the subsequent Lk with $k > 0$ will presumably monotonically move from bidding one's own valuation θ_i to bid $\theta_i/2$, but it still has to be computed.

Actually, from back-of-the-envelope calculations it seems that the best reply against a bidder who bids in $U[0,1]$ or against a bidder who bids in $U[0,1/2]$ is always to bid $\theta_i/2$:

$$\arg\max_{x \leq \theta_i} x(\theta_i - x) = \arg\max_{x \leq \min\{\theta_i, 1/2\}} 2x(\theta_i - x) = \frac{\theta_i}{2}$$

Therefore, L1 already coincides with the BNE. The same is true in case of n bidders, where the L0 strategy is to bid θ_i, and the L1 strategy coincides with the BNE, namely, to bid $\frac{n-1}{n}\theta_i$.

4.6 Different information partitions across players

So far we only considered cases when the distribution F is partitioned equally across player. How about the case in which information is partitioned differently across players?[35] Our non-equilibrium concept can cope with this case too, in a simple way that is readily understood if we just apply our non-equilibrium to the game proposed in Brocas et al. (2014).

There are three equally likely states of the nature: A, B, and C. Player 1 privately observes whether the state of nature is A or {B,C}. Player 2 privately observes whether the state of nature is C or {A,B}. Thus, the information is partitioned differently across players. Players choose to either bet on the state of nature, or not to bet and secure the sure payoff S (see the picture below, which is Figure 1 in Brocas et al., 2014). Payoffs are S unless both players decide to bet, in which case the payoffs are represented in the 2x3 matrix below, according to player and true state of nature.

	A	B	C	S
1	25	5	20	**10**
2	**0**	30	5	

F: 2 in $\{A\}$

Consider the behavior of player 1 who observes information {A,B}.[36] Then:

- If she is L0, she will ignore the true distribution of types, and she believes that her rival also observed state {A,B}. Thus, her rival (player 2) will bet

[34] More formally, when the valuation is θ for both bidders, every pair (x, θ) with $x \leq \theta$ is a NE of the complete information game. Otherwise, any winning bid greater than θ is irrational, and a winning bid lower than θ would yield incentives to overbid.

[35] We thank Vincent Crawford for suggesting to analyze how our non-equilibrium concept copes with this possibility.

[36] Notice that we slightly modify the original Brocas et al. game in order to avoid dominant strategy that implies a constant action at all levels of sophistication (like action S in their Figure 1 for a player 2 who observes state {A}).

because $(0+30)/2 > 10$, and this implies that player 1 will also bet because $(25+5)/2 > 10$.

- If she is L1, she correctly believes that her rival observes state {A} with probability 0.5 and observe state {B,C} with the remaining probability.[37] Also, she believes that her rival is L0, so that if her rival observes {A}, she does not bid, because $0 < 10$, while if her rival observes {B,C}, she does bid, because $(30+5)/2 > 10$. Therefore, back to the L1 player 1, upon observing {A,B}, she believes that if the true state is A, she will get 10, and if the true state is B, she will get 5. Thus, not bidding and securing 10 is the L1 choice.

- If she is Lk with k$\geq$ 2, the following reasoning holds. She will never receive the payoff 25 because in state {A} her rival, whatever her sophistication is, will play her dominant strategy S. Therefore, since any convex combination between 10 and 5 is lower than the payoff by choosing S (i.e., 10), any Lk with $k \geq 2$ will not bet. Note that not betting is the NE of the game, which is already achieved with one step of reasoning, in this case.

5 Formal definition of the non-equilibrium concept

In this section we adapt the standard definition of BNE to incorporate our level-k reasoning.

A Bayesian game consists of:

1. A set of players $\mathcal{N}$
2. A set of actions (pure strategies) for each player $i \in \mathcal{N}$: S_i
3. A set of types for each player i: $\theta_i \in \Theta_i$
4. A payoff function for each player i: $u_i(s_1, ..., s_n, \theta_1, ..., \theta_n)$
5. A (joint) probability distribution $p(\theta_1, ..., \theta_n)$ over types (or $P(\theta_1, ..., \theta_n)$ when types are not finite).

More generally, one could allow for a signal for each player, which is correlated with underlying vector of types. We do not consider this for the time being.

A pure strategy is a mapping s_i: $\Theta_i \to S_i$, selecting an action for each possible type of player i. Player i knows her own type and evaluates her expected payoff according to the conditional distribution $p(\theta_{-i} \mid \theta_i)$, computed by Bayes rule. In Bayesian games everything is common-knowledge. This is a strong but convenient assumption, since in a private information setting it allows players to form beliefs

[37] Notice that by definition of L1, she knows that the true state is either A or B. Thus, the probabilities of each state that her rival might be observing are computed by Bayes rule.

about types, and to understand others' beliefs about her own type, and so on. Thus, we can compute expected payoffs of player i of type θ_i as

$$U(s_i', s_{-i}(\cdot), \theta_i) = \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

then a Bayes equilibrium $s(\cdot)$ is defined as

$$s_i(\theta_i) \in \arg\max_{s_i' \in S_i} \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

If the Bayesian game is finite, then a mixed strategy exists. Existence of a pure strategy is guaranteed by continuous strategy spaces and continuous types, compact strategy sets and compact type sets, payoff functions continuous and concave in own strategies.

A (type-symmetric) L0 strategy solves

$$s_i^{L0}(\theta_i) \in \arg\max_{s_i' \in S_i} u_i(s_i', s_{-i}(\theta_i), \theta_i, ..., \theta_i)$$

in words, an L0-player of a certain type θ_i believes that all other players are of the same type θ_i and that this is commonly known.

Recursively, an Lk strategy with $k \geq 1$ solves

$$s_i^{Lk}(\theta_i) \in \arg\max_{s_i' \in S_i} \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}^{Lk-1}(\theta_i), \theta_i, \theta_{-i})$$

so that an L1-player best replies to the correct distribution of types, but believes that the other players are L0, and so on.

We conjecture that the existence of an Lk action $\forall k \in \mathbb{N}_+$ in "well-behaved games" should follow from the existence of an equilibrium in the symmetric complete information version of the game (i.e., L0 action), which seems a mild assumption to impose.

6 Conclusions

The literature on limited depth of reasoning in games with complete information is burgeoning, and yet considerably less is studied when the game has some private information. This paper provides a step forward in this direction, proposing an extension of the standard level-k reasoning to private information games. Information processing is at the heart of private information games. We propose to disentangle the infinite iteration of reciprocal expectations of players that lies behind the concept of a Bayes Nash equilibrium and propose a non-equilibrium concept where a player is characterized by her depth of reasoning k, which specifies the order of reciprocal expectation that she correctly computes. The anchoring behavior of players who do not compute any correct reciprocal expectation is to believe that their rival is of their same type. This "gut reaction" thus comes before the cognitive understanding of how to process the implications of the role of information in the game, and subsequently include such reasoning in the action's choice.

We propose a simple game to capture in the lab the level k that our non-equilibrium concept dictates, and we apply our non-equilibrium concept to standard games.

At least one key question is still open. In a standard p-beauty contest it is hard to disentangle when we observe that a subject picks a high number whether it is the result of a low depth of reasoning or of a belief of high depth of reasoning of others. Likewise, our Up-or-Down game suffers from the same identification problem. This problem triggered several strands of literature based on different approaches. One approach is to elicit subjects' beliefs; however, in theory one would have to elicit all orders of beliefs in order to correctly solve the issue; e.g., Costa-Gomes and Weizsäcker (2008) and Healy (2011) elicits first-order beliefs. An alternative approach is to design an experiment for which beliefs on others can be recovered directly from subjects' choices, such as the ring-network game proposed by Kneeland (2012). We hope to trigger efforts to solve this identification issue in a private-information setting too.

References

- Alaoui, L. and Penta, A., (2016a). "Endogenous Depth of Reasoning," *Review of Economic Studies*, Vol 83, Issue 4: 1297-1333

- Alaoui, L. and Penta, A., (2016b). "Cost-Benefit Analysis in Reasoning," Working Paper.

- Arad, A. and Rubinstein, A. (2012). "The 11-20 Money Request Game: A Level-k Reasoning Study," *American Economic Review*, 102 (7), 3561-3573.

- Alighieri, Dante. Paradiso canto XIII: 118–120. Trans. Allen Mandelbaum

- Anderson, C. A., Lepper, M. R. and Lee, R. (1980). "Perseverance of Social Theories: The Role of Explanation in the Persistence of Discredited Information," *Journal of Personality and Social Psychology*, 39 (6): 1037–1049.

- Baye, M. R., Kovenock, D. and de Vries, C. G. (1996). "The all-pay auction with complete information," *Economic Theory*, Volume 8, Issue 2, pp 291–305

- Biais, B., and Weber, M. (2009). "Hindsight bias, risk perception, and investment performance," *Management Science*, 55, 1018–1029.

- Blank, H., Musch, J., and Pohl, R. F. (2007). "Hindsight Bias: On Being Wise After the Event," *Social Cognition*, 25(1), 1–9.

- Brocas, I., Camerer, C., Carrillo, J. and Wang, S. (2014). "Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games," *Review of Economic Studies*, 81, 944-970.

- Buchanan, J. M. (1975) "The Samaritan's dilemma," in E. S. Phelps, ed., 'Altruism, morality and economic theory', Russell Sage Foundation, pp. 71–85.

- Camerer, C., HO, T.-H. and Chong, J.-K. (2004). "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119 (3), 861–898.

- Costa-Gomes, M., A., Crawford, V. P. and Broseta, B. (2001). "Cognition and Behavior inNormal-Form Games: An Experimental Study," *Econometrica*, 69 (5), 1193–1235.

- Crawford, V. P., and Iriberri, N. (2007). "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, Vol. 75, No. 6, pp. 1721-1770

- Costa-Gomes, M. and Weizsäcker, G. (2008) "Stated Beliefs and Play in Normal-Form Games," *Review of Economic Studies*, vol. 75, p. 729-762

- Fischhoff, B. (1975) "Hindsight / foresight: The Effect of Outcome Knowledge On Judgement Under Uncertainty", *Journal of Experimental Psychology*: Human Perception and Performance, Vol. 1, 288-299.

- Hargreaves-Heap, S., Rojo-Arjona, D. and Sugden, R., (2014) "How Portable is Level-0 behavior? A Test of Level-k Theory in Games with Non-Neutral Frames," *Econometrica* 82, 3, pp.1133-1151.

- Harsanyi, J. C. (1967–1968). "Games on incomplete information played by Bayesian players. Parts I–III." *Management Science* 14, 159–182, 320–334, 486–502.

- Healy, P. J. (2013), "Epistemic Foundations for the Failure of Nash Equilibrium," Working Paper.

- Kahneman, D. and Tversky, A. (1973). "On the psychology of prediction," *Psychological Review*. 80 (4): 237–251.

- Kahneman, D., Slavic, P. and Tversky, A., (1982) "Judgement under uncertainty: Heuristics and biases," Cambridge University Press, Cambridge

- Kneeland, T. (2012) "Identifying Higher-order Rationality," *Econometrica*, Vol. 83, 2065-2079

- Loewenstein, G., Moore D. and Weber, R. (2006) "Misperceiving the Value of Information in Predicting the Performance of Others", *Experimental Economics*, Vol. 9, 281-295.

- Lord, C. G., Lee, R., and Lepper, M. R. (1979) "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence," *Journal of Personality and Social Psychology*, 37 (11): 2098–2109,

- Madarasz, K. (2012) "Information projection: Model and applications," *Review of Economic Studies*, 79(3):961–985.

- Mueller, J.E., (1994). "Policy and Opinion in the Gulf War," University of Chicago Press, Chicago.

- Mullen, B., Atkins, J., Champion, D., Edwards, C., Hardy, D., Story, J., and Vanderklok, M. (1985). "The false consensus effect: A meta-analysis of 115 hypothesistests," *Journal of Experimental Social Psychology*, 21, 262–283

- Nickerson, R. S. (1998). "Confirmation bias: A ubiquitous phenomenon in many guises," *Review of General Psychology*, Vol 2(2) 175-220.

- Rogers, B. W., Palfrey, T. R., and Camerer, C. F. (2009). "Heterogeneous quantal response equilibrium and cognitive hierarchies", *Journal of Economic Theory*, Volume 144, Issue 4, Pages 1440–1467

- Ross, L., Greene, D. and House, P. (1977) "The “false consensus effect”: An egocentric bias in social perception and attribution processes", *Journal of Experimental Social Psychology* Volume 13, Issue 3, Pages 279–301

- Ross, L., and Anderson, C. A. (1982) "Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments", in Kahneman, Daniel; Slovic, Paul; Tversky, Amos, Judgment under uncertainty: Heuristics and biases, Cambridge University Press, pp. 129–152

- Rubinstein, A. (1989) "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"," *American Economic Review*, Vol. 79, No. 3, pp. 385-391

- Rubinstein, A. and Salant, Y. (2016) "“Isn't everyone like me?”: On the presence of self-similarity in strategic interactions", *Judgment and Decision Making*, Vol. 11, No. 2, pp. 168-173

- Schuett, F., and Wagner, A. K. (2011). "Hindsight-biased evaluation of political decision makers," *Journal of Public Economics*, 95, 1621–1634.

- Schelling, T. C. (1960). "The strategy of conflict," First ed., Cambridge: Harvard University Press.

- Shimp, T. A., and Sharma, S. (1987). "Consumer ethnocentrism: Construction and validation of the CETSCALE," *Journal of Marketing Research*, 24(8), 280-289.

- Sims, C. A. (2003). “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50, 665–90.

- Stahl, D. O. and Wilson, P. (1994) "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, 25 (3), 309–327.

- Wolfstetter, E. G. (1999). "Topics in Microeconomics", Cambridge University Press.