



**COLLANA DEL
DIPARTIMENTO DI ECONOMIA**

OBJECT-ORIENTED BAYESIAN NETWORKS FOR A DECISION SUPPORT SYSTEM

Julia Mortera - Paola Vicard - Cecilia Vergari

Working Paper n° 144, 2012

I Working Papers del Dipartimento di Economia svolgono la funzione di divulgare tempestivamente, in forma definitiva o provvisoria, i risultati di ricerche scientifiche originali. La loro pubblicazione è soggetta all'approvazione del Comitato Scientifico.

Per ciascuna pubblicazione vengono soddisfatti gli obblighi previsti dall'art. 1 del D.L.L. 31.8.1945, n. 660 e successive modifiche.

Copie della presente pubblicazione possono essere richieste alla Redazione.

REDAZIONE:

Dipartimento di Economia
Università degli Studi Roma Tre
Via Silvio D'Amico, 77 - 00145 Roma
Tel. 0039-06-57335655 fax 0039-06-57335771
E-mail: dip_eco@uniroma3.it



DIPARTIMENTO DI ECONOMIA

OBJECT-ORIENTED BAYESIAN NETWORKS FOR A DECISION SUPPORT SYSTEM

Julia Mortera - Paola Vicard - Cecilia Vergari

Comitato Scientifico:

Fabrizio De Filippis

Anna Giunta

Paolo Lazzara

Loretta Mastroeni

Silvia Terzi

Object-Oriented Bayesian Networks for a Decision Support System *

Julia Mortera

Paola Vicard

Università Roma Tre, Italy

Università Roma Tre, Italy

Cecilia Vergari

Università di Bologna, Italy

Abstract

We study an economic decision problem where the actors are two firms and the Antitrust Authority whose main task is to monitor and prevent firms potential anti-competitive behaviour. The Antitrust Authority's decision process is modelled using a Bayesian network whose relational structure and parameters are estimated from data provided by the Authority itself. Several economic variables influencing this decision process are included in the model. We analyse how monitoring by the Antitrust Authority affects firms cooperation strategies. These are modelled as a repeated prisoners dilemma using object-oriented Bayesian networks, thus enabling integration of firms decision process and external market information.

JEL Classification: C440, C730, L400, C110, D810, D830, L130.

Keywords: Antitrust Authority, Bayesian networks, mergers, model integration, prisoners dilemma, repeated games.

*Sections 2.1, 2.3 and 2.4.2 are due to C. Vergari the remaining sections are due to J. Mortera and P. Vicard.

1 Introduction

Firms in many cases have incentives to cooperate (collude) to increase their profits. The possibility of firms to collude does not depend solely on their decision but also on external circumstances. First of all, firms need to comply with antitrust laws. If the AA finds negative anti-competitive effects, resulting from firms cooperative behaviour, it may intervene to prevent the firms from merging. The AA's decision process is modelled here by using a Bayesian network (BN) or Probabilistic Expert System (PES) (Cowell *et al.* 1999) estimated from real data. A BN is a graphical model that encodes the probabilistic relationships among the variables of interest allowing for the application of fast general-purpose algorithms to compute inferences.

We also study how the AA monitoring affects firms' strategies about cooperation. To this aim, firms set of potential strategies are modelled in turn as a repeated prisoner's dilemma using object-oriented Bayesian networks (OOBNs) (Koller and Milch 2001; Bangsø and Wuillemin 2000). OOBNs are a recent extension of BNs which allow for a hierarchical definition and construction of a BN. They provide a compact and intuitive representation of the repeated prisoner's dilemma. Furthermore, thanks to the modularity and flexibility of this approach, various sources of uncertainty within the game and generalizations of the repeated prisoner's dilemma can be analysed.

Here we present two different networks: the first to represent the duopoly, and the second to model the AA's decision process. OOBNs give the graphical framework to integrate these two networks and to represent their time evolution. Both the graphical structure and the associated probability tables of AA's decision process network are estimated from a real dataset. As an outcome, we obtain the estimated probability that AA intervenes to prevent anticompetitive behaviour of a merger. In a general setting, this corresponds to the probability that there is no successive stage in a repeated prisoner dilemma (PD). For various economic sectors (markets of interest) we study the sensitivity of cooperative outcomes with respect to factors such as geographical size, market share, Herfindahl-Hirschman Index (HHI) variation, vertical effects, the presence of entry barriers and buyer power. The global OOBN model which integrates the AA's decision process with a duopoly

model is used to obtain the optimal decision in light of a series of interesting scenarios that could occur in practice.

The outline of the paper is as follows. We first give a brief introduction to the prisoners dilemma in Section 2.1 followed by the Bayesian network representation of the PD in Section 2.2. After introducing the repeated prisoners dilemma in Section 2.3, in Section 2.4 we show how this can be represented as an OOBN. In Section 3.1 we illustrate the BN for the AA's decision process estimated from the data and show its use in various scenarios. In Section 3.2 we show how we integrate the PD network with the AA network obtaining a general purpose global representation of the problem, and in Section 3.3 we apply this to several decision scenarios. Finally, in Section 4 we draw conclusions and discuss further developments.

2 DUPOPOLY REPRESENTATION

2.1 The prisoner's dilemma

The prisoner's dilemma (Flood 1958) describes cooperation by rational agents. The PD is a 2-player symmetric game where the two players have the same rôle and have the same set of potential strategies termed *cooperate* C , and *defect* D . The PD is a simultaneous game where the players choose just once and simultaneously and the unique equilibrium¹ is the pair of strategies (D, D) . Players' payoffs are such that defect is a dominant strategy, *i.e.* a strategy that is preferred by each player independently of his/her rival. The problem is that this strategy is inefficient since both players would gain more if they cooperated and adopted the (C, C) strategy. The source of the dilemma lies in the fact that each player has an incentive to defect if the rival player cooperates; so that an agreement to cooperate would not be credible.

Simultaneous games, such as the PD, are commonly represented in either the normal or the extensive form. In the normal form representation, the PD can be described by the payoff matrix in Table 1. The two firms, Firm1 and

¹An equilibrium is a strategy pair such that no player can improve his position by unilaterally changing his decision. In other words, it is a situation in which all players choose mutual best responses.

Firm2, have two available strategies: cooperate C , or defect D . The payoffs need to be such that $d > a > b \geq c$ and $2a > (c + d) > 2b$, so that (C, C) maximizes players' joint payoff. Given that $b < a$, the strategy pair (D, D) is strictly worse than (C, C) .

		Firm2	
		C	D
Firm1	C	a,a	c,d
	D	d,c	b,b

Table 1: Payoff matrix for the prisoner's dilemma

In the extensive form the game is represented by a tree. Figure 1a shows the tree representation (equivalent to Table 1) of the simultaneous duopoly game. Firm1 moves first and chooses either C or D , Firm2 moves second but without knowing what Firm1 did.




Figure 1: a) Tree representation of the simultaneous duopoly game. b) Corresponding Bayesian network representation.

A symmetric duopoly model, such as a market with two symmetric profit-maximising firms in mutual competition, is an example of a PD. The duopoly profit is the gain of each of the sellers in this market.

Suppose the two firms produce identical goods, incurring constant marginal costs, and they compete setting their prices. Since consumers will buy from

the firm charging the lowest price, firms have an incentive to undercut their price to conquer the market (non-cooperative or defect strategy). At equilibrium firms will set the competitive price (the market price under perfect competition which is equal to firms marginal cost of production) gaining duopoly profit $b = 0$. This result is often called a paradox, since there are just two firms in the market and still the perfectly competitive strategy yields zero profit. However, if firms decide to cooperate and set the monopoly price, they can share positive monopoly profits. The monopoly profit is always greater than twice the duopoly profit, $2a > 2b$.

In most markets, from a consumer's point of view goods are not identical. This gives firms the ability to raise the price above the marginal cost of production without losing their customers to competitors. In a symmetric duopoly with product differentiation firms produce and sell differentiated goods (imperfect substitutes). As long as product differentiation is not too large, firms face a PD: if they cooperate they could share monopoly profit, but they have incentive to defect if the rival cooperates. However, when goods are imperfect substitutes, firms make positive duopoly profit, $b > 0$, under the non-cooperative strategy pair (D, D) . This duopoly profit is smaller than half the monopoly profit, $b < a$, so that the cooperative strategy C is superior for each firm singly.

2.2 The prisoners dilemma network

Bayesian networks for decision support systems can incorporate both decision nodes and utility nodes (Jensen 2001) giving rise to an influence diagram (ID) representation. IDs were extended by Lauritzen and Nilsson (2001) to allow for limited information decision problems (LIMIDs). A different approach to represent and solve games using graphical models was initially proposed by Smith (1996), and later by La Mura (2000), Kearns *et al.* (2001) and Koller and Milch (2003).

The one stage PD being a symmetric game can be represented by the ID network in Figure 1b. The simultaneity of the game is implemented by representing Firm1 as a random variable (oval node), and Firm2 as the decision maker (rectangular node) having two possible actions: defect D , and coop-

erate C . Firm2's decision is influenced by Firm1. Firm1's associated prior probability distribution represents Firm2's subjective opinion about Firm1's behaviour. Random variable Firm1 has two states, defect (coded as 0) and cooperate (coded as 1), with uniform prior probabilities indicating Firm2's ignorance about Firm1's choice. Firm2 could assign different prior probabilities based on his/her prior knowledge about Firm1's behaviour. Table 2 shows Firm2's utility (node Firm2's utility U2 in Figure 1b) based on Firm1 and Firm2's actions. Thanks to game symmetry, Table 2 is equivalent to the normal form payoff matrix given in Table 1.

Table 2: Firm2's utility U2 conditional on Firm1 and Firm2's actions.

Firm1	defect (0)		cooperate (1)	
Firm2	defect (0)	cooperate (1)	defect (0)	cooperate (1)
U2	b	c	d	a

Once the network is compiled, the optimal decision for Firm2 is automatically computed by maximising expected utility. Since the game is symmetric, Firm2's optimal strategy coincides with Firm1's optimal strategy and this pair of strategies constitutes a Nash equilibrium. Thus in the ID representation the choice of Firm2 as decision maker is without loss of generality.

In what follows we always consider Firm2 as the decision maker. The prior probability distribution on the random variable Firm1 reflects Firm2's subjective opinion on the type of rival player he/she is playing against. Kadane and Larkey (1982) explore the consequences for game theory of adopting a subjective view of probability. They show that assumptions like, $n = 2$ versus $n > 2$ person game, zero sum versus variable sum, are not critical; whereas, the distinction between single and repeated game is essential in this framework.

2.3 Repeated prisoner's dilemma

In repeated games, players actions are observed at the end of each period and their overall payoff is the sum of the payoffs in each stage discounted by

a factor $\delta \in [0, 1]$. Thus players may condition their play on the opponents past play. Here we assume that firms never forget previous moves and other information acquired, in other words we assume that firms have perfect recall. This can lead to equilibrium outcomes that differ from simultaneous game outcomes (the folk theorem, first proved by Friedman (1971), provides a formal proof of this result).

The repeated PD analyzes how threats and promises about future behaviour can affect and improve current behaviour. A key ingredient for cooperation to arise among rational players is to have an infinite or indefinite horizon. An infinitely repeated game is equivalent to a game repeated an indeterminate number of times, *i.e.* a game where each future stage takes place with probability δ (Roth and Murnighan 1978). Thus δ represents uncertainty about the number of stages faced by firms. This uncertainty is usually not modelled within the game itself.

In a duopoly, where repeated interactions between firms take place, cooperation can arise as an equilibrium of the repeated game. Consider for example a duopoly game repeated an infinite number of times. Suppose that Firm1 adopts the strategy (named grim trigger strategy) to cooperate as long as Firm2 cooperates and, if in any period Firm2 defects, then Firm1 defects in every subsequent stage. If Firm2 cooperates in each stage, both firms gain the cooperative profit a . Instead, if in any stage, Firm2 defects it gains the deviation profit, $d > a$, in that stage and the non-cooperative duopoly profit, $b < a$, in all successive stages. In a duopoly, δ represents an evaluation of firms' future payoffs. Thus Firm2 prefers to cooperate if δ is sufficiently large as the short-run gain from defection $d - a$ are overcompensated by the long-run losses $b - a$ in each successive stage.

Generalizing the tree representation in Fig 1a to repeated games is both computationally and graphically demanding. The game tree grows exponentially with the number of stages. For example, Figure 2a shows the tree representation of a two-stage PD.

In general, for an n -stage game there are 4^n outcomes (tree leaves). Furthermore, one needs to add another layer for each stage to represent uncertainty about the occurrence of a future stage.

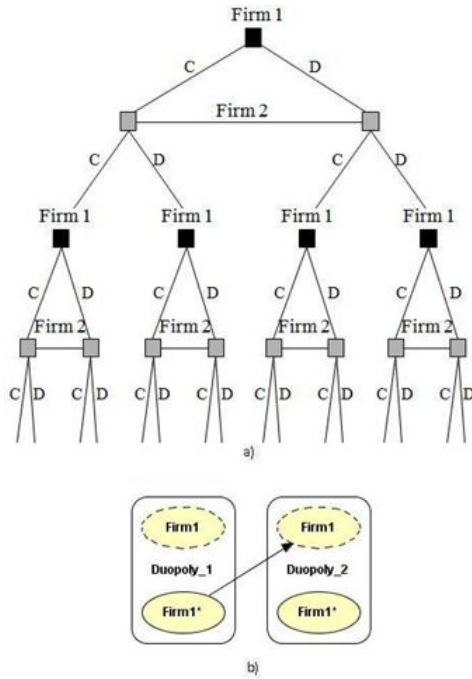


Figure 2: a) Tree representation of the two-stage duopoly game. b) Corresponding OOBN representation.

2.4 OOBN for repeated prisoner's dilemma

A potentially promising new approach for dealing with symmetric repeated games is to reformulate them as an OOBN. Object-oriented Bayesian networks have a hierarchical structure where a node itself can represent a (object-oriented) network containing several *instances* of other generic *classes* of networks. Instances have interface *input* and *output* nodes as well as ordinary nodes. Instances of a particular class have identical conditional probability tables for non-input nodes. Instances are connected by arrows from output nodes into input nodes. These arrows, as well as those from ordinary nodes to input nodes, represent identity links, whereas arrows between two ordinary nodes or an output node and an ordinary node represent probabilistic dependence.

OOBNs are particularly well suited for an application area such as the present because the similarity between network elements (the stages of the

game) can be exploited in a modular and flexible construction. The OOBN is not only a pictorial representation of the game but it also incorporates an inference engine which is also the game solver. The graphical simplicity automatically produces computational efficiency. As a result, increasingly complex networks can be constructed by simply adding new objects which perform different tasks. Figure 2b shows the OOBN two-stage repeated game that corresponds to the tree representation in Figure 2a. Compared to the tree representation in Figure 2a, Figure 2b has a remarkably simpler graphical representation.

Since we assume perfect recall, HUGIN² version 6.9 software, which automatically implements the fact that at every stage the decision maker recalls all previous decision, is used to build the networks. In what follows we indicate an instance in **bold** and a node in **teletype**. In the OOBN of Figure 2b each rounded rectangle represents an instance termed **Duopoly** and models a stage of the repeated game. In order to specify the links between successive stages (instances) Figure 1b (which represents each **Duopoly** instance) needs to be generalized as shown in Figure 3.

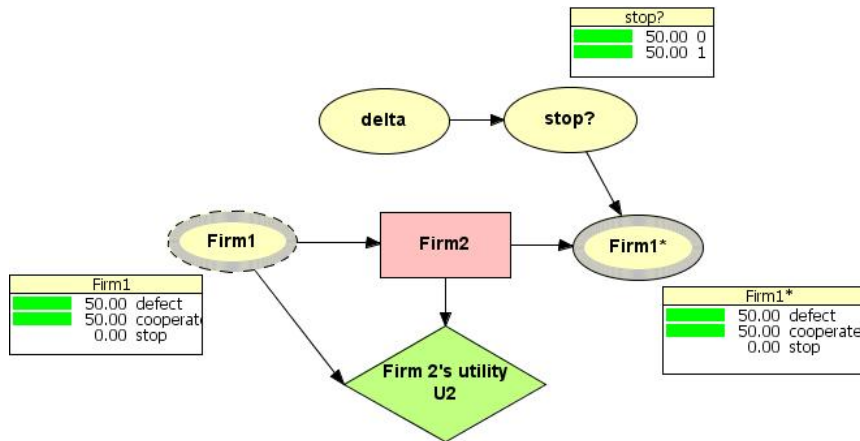


Figure 3: Class network for repeated PD with associated marginal prior probability tables.

The node **Firm1*** models the behaviour of Firm1 in the next stage. In

²www.hugin.com

each stage the game can either continue or terminate. `Firm1` and `Firm1*` now need to be given three states: defect (0), cooperate (1) and stop (2). Since in a repeated game every stage depends on the actions taken in the previous stages, `Firm1*` is logically dependent on `Firm2`. Uncertainty about the existence of further stages is modelled by adding a new random node `stop?`. Node `stop?` has two states, $\{0, 1\}$ according to whether the game continues or stops and has a Bernoulli distribution $Bin(1, 1 - \text{delta})$. The parameter node `delta` is the probability that the game continues $P(\text{stop?} = 0)$. It represents the discount factor δ described in Section 2.3. Node `delta` has a uniform prior distribution over a plausible set of values.

In the first stage to ensure that the game starts `Firm1` can only choose between defect and cooperate. Table 3 gives the conditional probability distribution of `Firm1*` given `stop?` and `Firm2`. It shows that if the game stops (`stop?=1`) `Firm1*` stops with certainty, else `Firm1*` cooperates or defects according to `Firm2`'s decision. This implements the *tit for tat* (TFT) strategy in which `Firm1` begins by cooperating and cooperates as long as `Firm2` cooperates, and defects otherwise. Variations on this strategy will be shown in Section 2.4.1.

Table 3: Conditional probability table for `Firm1*` given `stop?` and `Firm2`.

<code>stop?</code>	no (0)		yes (1)	
	defect (0)	cooperate (1)	defect (0)	cooperate (1)
<code>Firm2</code>				
defect (0)	1	0	0	0
cooperate (1)	0	1	0	0
stop (2)	0	0	1	1

The key issue to represent indefinitely repeated games is to model the uncertainty about the occurrence of the next stage, as represented by node `stop?`. In general, an OOBN with $n + 1$ instances models a game repeated n times with uncertainty about the successive stage.

2.4.1 Other Strategies

Bayesian networks can also be usefully generalised to model strategies other than TFT. Figure 3 can be modified to provide a general class network that incorporates a set of potential strategies for Firm1. This network is displayed in Figure 4. Additional nodes, $Firm1^*|D$ and $Firm1^*|C$, having Bernoulli distributions with parameter nodes $alpha_D$ and $alpha_C$ are added to the network of Figure 3. Node $Firm1^*$ takes value 2 if the game stops in the current stage, whereas, if the game continues ($stop?=0$), the value of $Firm1^*$ depends on that of $Firm2$. If $Firm2$ defects $Firm1^*$ is $Firm1^*|D$, else $Firm1^*|C$. The conditional probability distribution of $Firm1^*$ is thus defined by the logical expression $if(stop == 1, 2, if(Firm2 == 0, Firm1^*|D, Firm1^*|C))$.³ We can thus represent $Firm2$'s subjective opinions about $Firm1$'s behaviour in each single stage of the repeated game. The TFT strategy can also be implemented using this more general network as will be shown in Section 3.3.




Figure 4: Further refinement of the PD stage game.

2.4.2 Incomplete information

An alternative interpretation of the OOBN in Figure 4 is a model for the repeated PD with *incomplete information*. This implies uncertainty about

³The function $if(A, x, y)$ takes value x if condition A is satisfied, otherwise y .

the type of rival that a firm is going to face. Actually, experimental results show that people, contrary to standard prescriptions of game theory, may cooperate more frequently than expected (Andreoni and Miller 1993). An explanation behind this empirical evidence is provided by the theoretical models of Kreps and Wilson (1982) and Kreps *et al.* (1982).

In our model Firm2's beliefs about its rival are modelled in Firm1*. The conditional probability distribution of Firm1* reflects Firm2's uncertainty about its opponent. If Firm2 believes Firm1 to be "altruistic" it can expect Firm1 to cooperate (with a probability $\alpha_D > 0$) even if it defected in the previous stage. On the other hand, if Firm2 believes Firm1 to be "egoistic", then it expects Firm1 to cooperate with probability $\alpha_C < 1$ even if it cooperated in the previous stage.

This model can also incorporate a large set of strategies and it can model scenarios where the probability the game continues depends on external factors. An illustrative example is given in Section 3.

3 Application to the Italian Antitrust Authority decision process

Often governments may find negative anti-competitive effects resulting from a merger. As a consequence, the decision by firms to cooperate is actually affected by the decision process of the AA. The AA may start an investigation either because two firms make a formal request to merge (explicit collusion) or because the authority suspects that two firms are implicitly colluding. In what follows the term merger will be used for both explicit and implicit collusion.

The AA studies the impact of a merger on the market and its consequences on social welfare. Hence, the AA's decision affects the future stages of the game as well as its equilibrium outcome. Thus, when choosing between cooperating or defecting, firms take the decision process of the AA into account. Firms generally consider the risk of an antitrust investigation in the case of implicit collusion.

In what follows, the actors are: the duopolists (Firm1 and Firm2) and

the Antitrust Authority. The global model, where both the duopoly and the AA's decision process are represented by OOBNs is shown in Figure 5. The **Duopoly** network (the bottom network) in Figure 5 is described in Section 2.4 and 2.4.1. It is similar to the network in Figure 4 except that the uncertainty about the next stage δ is now replaced by a complex network (the top network in Figure 5) termed **AA**. This network represents the decision process of the AA and is estimated from real data. We will first show how we derive the network for the AA decision process and then we illustrate the integrated network and its use.

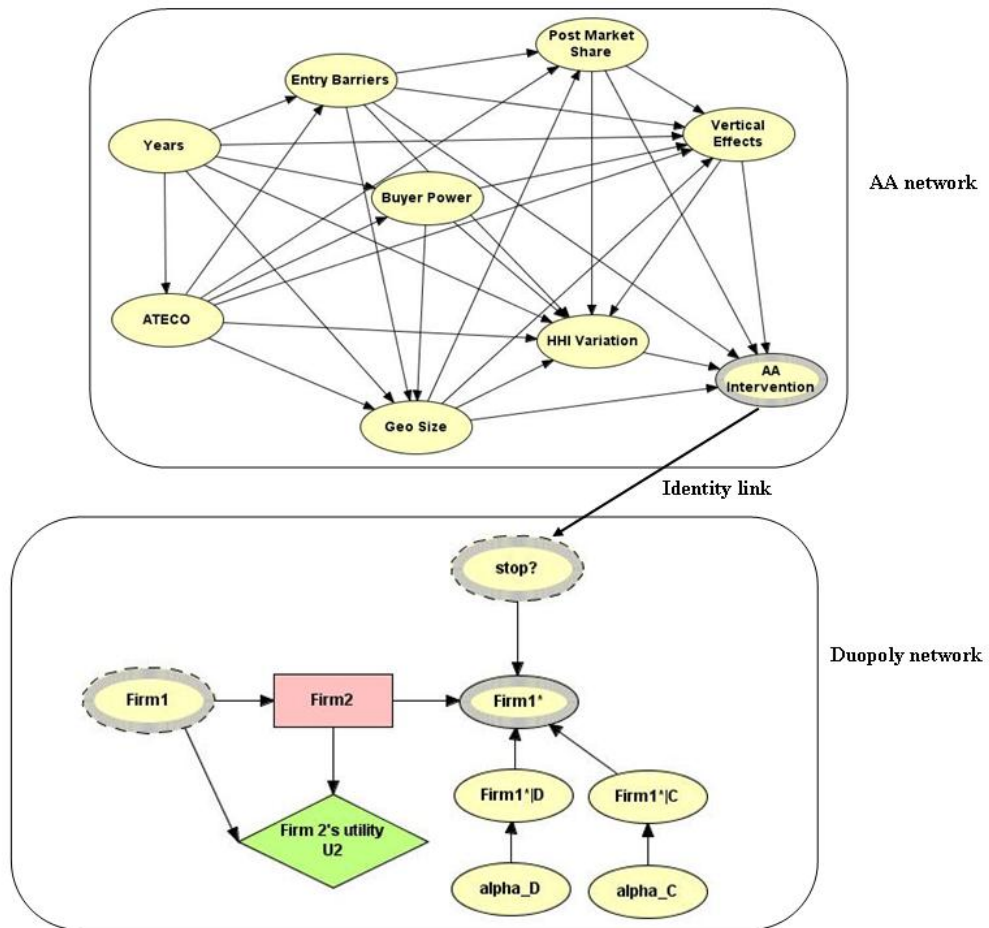


Figure 5: Integrated AA-duopoly merger stage game.

3.1 Antitrust Authority network

3.1.1 The data

The data we use were collected by the Italian Antitrust Authority and concern all the cases examined from 1991 to 2003. This dataset consists of 6920 observations. Based on this dataset, La Noce *et al.* (2006) developed a logit model to analyze the impact of different factors on the Authority decision. Following La Noce *et al.* (2006) we consider relevant markets affected by the merger as elementary units of analysis. These markets are denoted by the ISTAT (Italian National Institute of Statistics) economic activity code ATECO. Table 4 describes the variables in the dataset that were used to estimate the AA network.

Table 4: Description of the variables in the AA network.

Variable	States	Description
Years	{1991–1996, 1997–2000, 2001–2003}	Reference periods
ATECO	Mining, Food & Beverage Manufacture, <i>etc.</i> (see Fig. 7)	Relevant market
Geo Size	{Sub-national, National, Supra-national}	Size of the relevant market
Buyer Power	{Yes, No}	Presence (Yes) of competitive pressure on the merging parties
Entry Barriers	{Yes, No}	Presence (Yes) of entry barriers
HHI Variations	{0, (0,100), [100, 500), [500, 1000), ≥ 1000 }	Variation in market concentration index
Post Market Share	{<20%, [20% 40%], >40%}	Post-merger market share
Vertical Effects	{Yes, No}	Presence (Yes) of vertical effects
AA Intervention	{0, 1}	No (0)/Yes (1)

The estimation (learning) process of a Bayesian network consists of two phases: the graphical structure estimation and the conditional probability

table estimation. These will be illustrated in turn.

3.1.2 Estimation of the network’s graphical structure

The graphical structure of the AA network representing AA decision process is obtained by a combination of subject-matter knowledge, provided by a domain expert, and the information in the data.




Figure 6: Logical constraints imposed in learning AA network.

The *Necessary Path Condition* (NPC) algorithm (Steck 2001) implemented in HUGIN is used to estimate the graphical structure of the network. The NPC algorithm takes into account logical constraints, such as, presence/absence of a link or assignment/ban of a specific direction between variables. The logical constraints we implemented here are shown in Figure 6. These imply that if there is a relation between two variables in different boxes, it must have the same direction as that in Figure 6. Furthermore, if two variables belong to the same box, their association (if it exists) can be in any one of the two possible directions. For example, if node **AA Intervention** is connected with any of the other variables, the direction has to be from these into **AA Intervention** node (AA decision logically depends on the values of the other variables). This means that arrows from **AA Intervention** to any other variable are logically prohibited. The reference period (node **Years**) is not influenced by any of the other variables in the model.

The dependence structure — based on the logical constraints given in Figure 6 — learnt from the data is shown in Figure 7. The main dependence relationships estimated from the data are:

- i) The market of interest (**ATECO**) can depend on **Year**: an economic sector could be more relevant and worth investigating during one of the three




Figure 7: AA network showing the dependencies of AA on the relevant variables describing the market and the marginal probabilities of the variables.

reference periods (note that the president of the AA changed in 1997 and from 2001 Italian currency Lira was replaced by the Euro).

- ii) **AA Intervention** depends directly on **HHI Variation**, **Vertical Effects**, **Post Market Share**, **Geo Size**, and **Entry Barriers**. Furthermore, the relevant market (**ATECO**) does not affect **AA Intervention** directly but only through the relevant features of the market and of the merging firms (**HHI Variation**, **Vertical Effects**, **Post Market Share**, **Geo Size**, and **Entry Barriers**).
- iii) The Herfindahl-Hirschman concentration index variation (**HHI Variation**) depends on all the variables that logically precede it or are on an equal footing (as shown in Figure 6). Whereas **Post Market Share** depends only on **Entry Barriers**, **Geo Size** and **ATECO**. An explanation of this could be that when a market sector is characterized by entry barriers (because of patents or increasing returns to scale) we expect that this market may be composed of a few firms with high market shares, thus influencing **Post Market Share** and a relevant **HHI Variation**.

Many other conditional independencies can be read off the AA network in Figure 7, but for brevity they will not be presented here.

3.1.3 Estimation of the probability tables

To complete the construction of our model, we estimate the conditional probability distributions of the variables from the data. The EM-algorithm (Dempster *et al.* 1977) is used for learning the probabilities. In order to avoid that certain configurations in the conditional probability tables have zero probability we set the prior probabilities according to Buntine (1991).

Figure 7 displays the marginal probabilities estimated from our data. Note, for example, that the probability of an AA intervention is only 0.0189 which could be due to the fact that in most cases, 74.38%, the post market share is less than 20% and entry barriers and vertical effects are absent (with probability 0.9793 and 0.9268, respectively), HHI index is less than 100 in 87.85% of the cases and only in 15.38% the geographical size is supra-national.

3.1.4 Using the network

Once the model has been estimated, we can address a number of questions about the AA's decision process. Various possible scenarios can be examined by inserting and propagating the appropriate evidence throughout the network. We illustrate three hypothetical scenarios.

Scenario A. What is the probability of an AA intervention in a merger request when there are entry barriers in the market? This scenario is represented in Figure 8a. The posterior probability of an **AA Intervention** increases from 0.0189 to 0.5790 when the evidence **Entry Barriers = Yes** is inserted and propagated throughout the network.

Scenario B. How would the probability obtained in Scenario A change if the Herfindahl-Hirschman concentration index variation (**HHI variation**) is in the class [100, 500)? Note in Figure 8b that the probability of **AA Intervention** now increases to 0.7741.

The network can be used not only for direct reasoning about the probability of **AA Intervention**, but also for reasoning about possible "causes" of a given AA decision.

Scenario C. Suppose that the AA decides to intervene in a firm's merger request. What are the most plausible reasons of this decision? Figure 8c gives the posterior probabilities given the evidence that **AA Intervention** is equal to one. On comparing Figures 7 and 8c we see that:

- the probability of entry barriers increases from 0.0207 to 0.6367;
- the probability of vertical effects increases from 0.0732 to 0.4536;
- the probability of post market share less than 20% decreases from 0.7438 to 0.0922, whereas the probability of post market share greater than 40% increases from 0.0777 to 0.7006.
- The HHI index decreases in the first two classes and increases in the last three classes.

Note that when evidence is propagated in the network, all marginal probability tables are updated accordingly.



Figure 8: Scenarios a), b) and c) giving marginal posterior percentages for the AA network.

3.2 Global network

Thanks to the modularity and flexibility of OOBNs, it is possible to integrate the AA and the Duopoly networks, giving rise to a unique overall OOBN representation of the problem. An expanded representation of this model is shown in Figure 5. The **Duopoly** network is an instance of the class network in Figure 4 where uncertainty about future stages, node **stop?**, is identified with **AA Intervention** in the **AA** network.

The AA decision process is usually dynamic, it can change over time due to changes in the antitrust law as well as changes in market conditions. We are thus interested in the repeated version of the model in Figure 5.

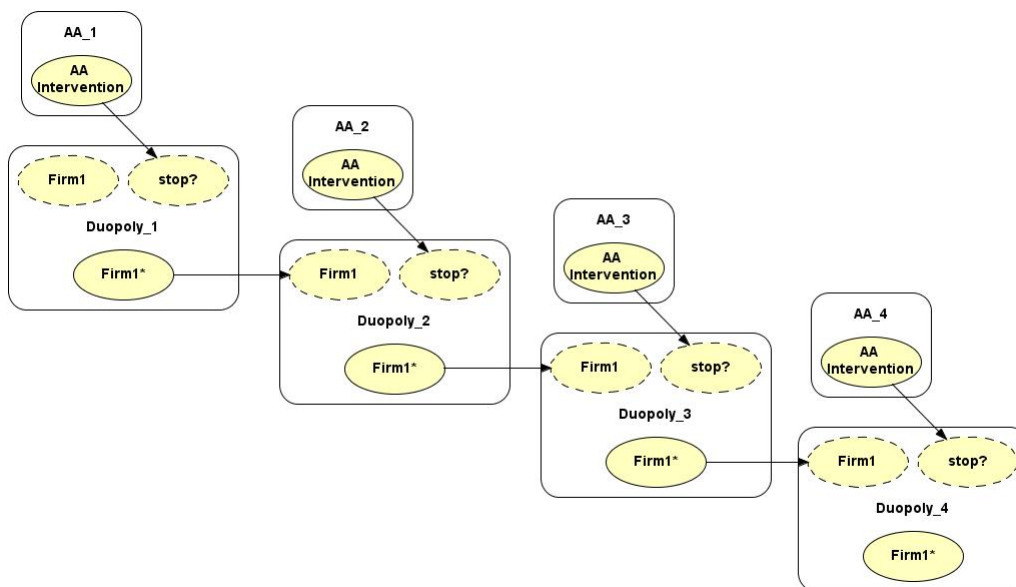


Figure 9: OOBN representing a three-stage repeated merger game with uncertainty about the number of stages.

Figure 9 represents the global model (Figure 5) repeated four times for a three stage merger game with uncertainty on the number of stages. In this model, the AA's decision process is represented by the same instance in each period. This is justified by assuming that even once the AA decides not to intervene, it continues monitoring firms' behaviour in successive stages.

3.3 Firms' strategy

We now study the sensitivity of cooperative behaviour with respect to two sets of utilities and all the factors that might directly or indirectly influence the AA's decision. We consider both the TFT strategy and a more general strategy. The TFT strategy can be implemented using the global network by setting $\text{Firm1}^* = 1$ in stage **Duopoly_1** and $\text{Firm1}^*|C = 1, \text{Firm1}^*|D = 0$ in all other stages.

3.3.1 TFT strategy: perfect substitutability

Table 5 shows an example of Firm2's utility for a market with perfect substitutable goods. Figures 10, 11 and 12 show the marginal probabilities for

Table 5: Firm2's utility U2 for a market with perfect substitutability

Firm1	defect (0)		cooperate (1)	
Firm2	defect (0)	cooperate (1)	defect (0)	cooperate (1)
U2	0	-10	150	100

a selection of random variables and the expected utilities for the decision nodes in the first stage **AA_1** and **Duopoly_1**.

When no evidence about the variables in the market is inserted in the network (Figure 10) Firm2's optimal decision is to cooperate (1) having expected utility equal to 443.40 (while defect has expected utility equal to 385.47). This could be in part due to the small probability of an AA intervention, 0.0189.

Figure 11 shows the case where there are entry barriers in the market of interest (**Entry Barriers = Yes**) and the merger causes the HHI variation to be in the last class (**HHI Variation ≥ 1000**). The resulting probability of AA intervention shoots up to 0.9435 and Firm2's optimal decision is to defect with expected utility of 394.72, against 350.93 for cooperating. This strategy still remains optimal (although with a smaller gap between the expected utilities) when based only on the presence of entry barriers.

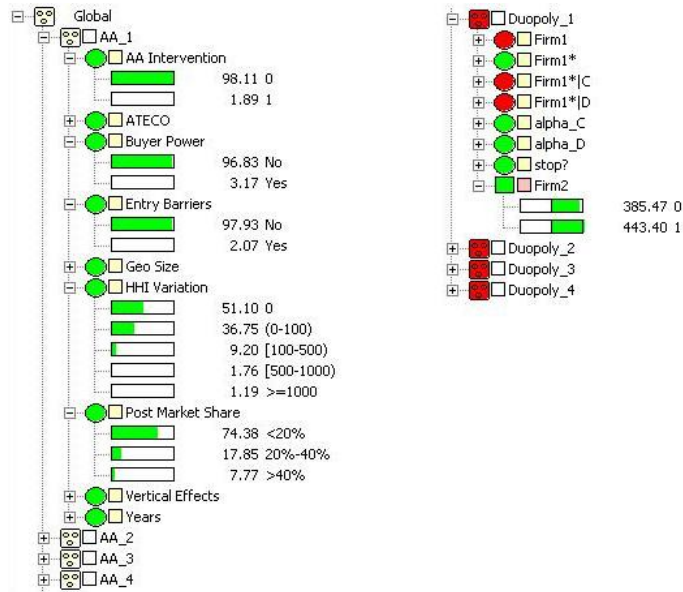


Figure 10: Marginal probabilities and optimal decision in the first stage **AA_1** and **Duopoly_1** when Firm1 plays TFT.

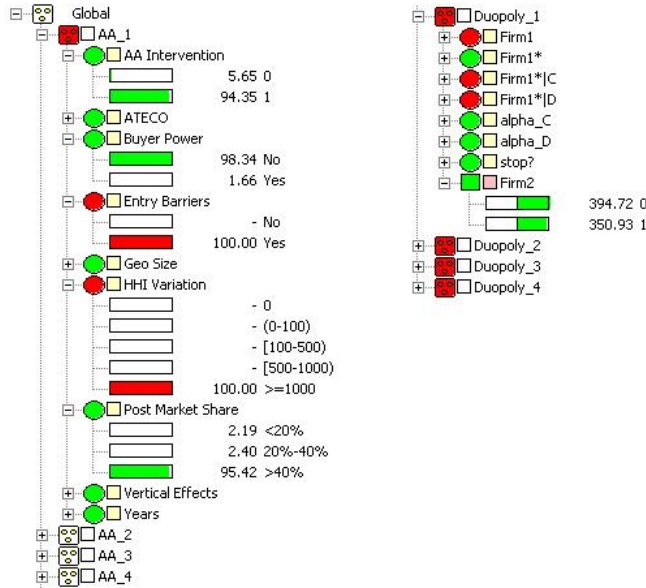


Figure 11: Marginal probabilities and optimal decision in the first stage **AA_1** and **Duopoly_1** when Firm1 plays TFT, **Entry Barriers** = Yes and **HHI Variation** ≥ 1000 .

Figure 12 shows the case where, as before, there are entry barriers, the HHI variation is ≥ 1000 , and customers exert competitive pressure on the merging parties (**Buyer Power** = Yes). The probability of AA intervention decreases from 0.9435 to 0.2915 and Firm2's optimal decision is to cooperate having expected utility of 416.14. It is interesting to note that buyer power is able to counterbalance the effect of both entry barriers and a large HHI variation.




Figure 12: Marginal probabilities and optimal decision in the first stage **AA_1** and **Duopoly_1** when Firm1 plays TFT, **Entry Barriers** = Yes, **HHI Variation** ≥ 1000 and **Buyer Power** = Yes.

3.3.2 TFT strategy: imperfect substitutability

We now use Firm2's utility for a market with imperfect substitutability given in Table 6. Figure 13 shows results when evidence about the market is not available. Firm2's optimal decision is to cooperate (1) having expected utility equal to 601.33 (while defect has expected utility equal to 513.21). Again, this is most plausibly due to the small probability of an AA intervention.

Table 6: Firm2's utility U_2 for a market with imperfect substitutability.

Firm1	defect		cooperate	
	defect	cooperate	defect	cooperate
Firm2	100	50	160	150

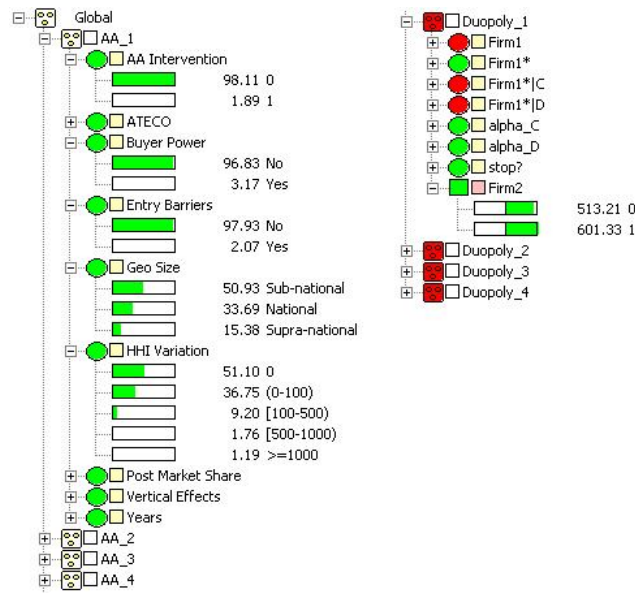


Figure 13: Marginal probabilities and optimal decision in the first stage **AA_1** and **Duopoly_1** under imperfect substitutability.

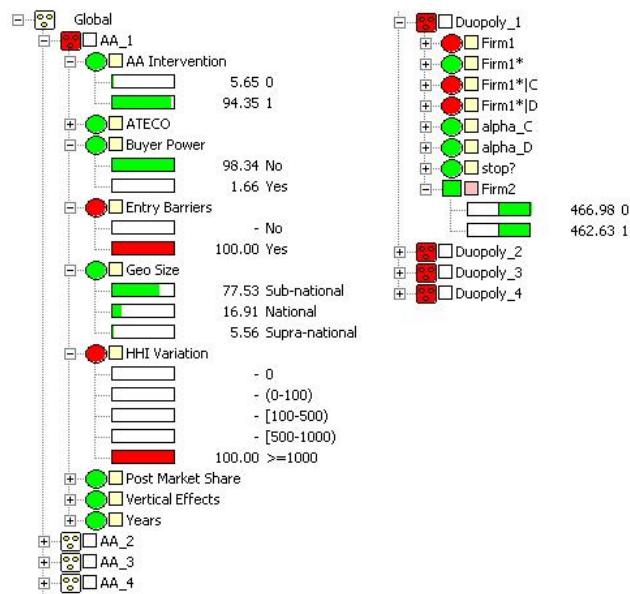


Figure 14: Marginal probabilities and optimal decision in the first stage **AA_1** and **Duopoly_1** under imperfect substitutability, when **Entry Barriers** = Yes, **HHI Variation** ≥ 1000 and **Buyer Power** = Yes.

Firm2's expected utility to cooperate or to defect are almost equal, although the probability of AA intervention is close to 1 (Figure 14) when `Entry Barriers = Yes` and `HHI Variation >= 1000`.

Furthermore, in contrast to perfect substitutability, accounting for the presence of entry barriers alone is not sufficient to modify the optimal decision from cooperate to defect. The main reason being that when the firms' products are imperfect substitutes, the set of utilities reflects the fact that the defect strategy does not correspond to such a strong punishment, so that a firm can continue to cooperate even if there is high risk that the game might stop.

3.3.3 Incomplete information

Assume that Firm2 has incomplete information about the type of rival it is going to face. This is a reasonable scenario as firms are likely to be uncertain about their rivals' costs and benefits from cooperation.

Table 7 shows Firm2's expected utility in case of perfect substitutability (based on Firm2's utility given in Table 5) for different probability values of α_C and α_D (nodes `alpha_C` and `alpha_D` in Figure 5). Three types of information about the relevant market are considered: no evidence, evidence $E_1 = \{\text{Post Market Share} \geq 40\%, \text{Entry Barriers} = \text{Yes and Buyer Power} = \text{Yes}\}$ and evidence $E_2 = \{\text{Entry Barriers} = \text{Yes and HHI Variation} \in [500-1000]\}$. The optimal decision yielding the highest expected utility for each scenario is italicised.

The second last row of Table 7 gives the results when inserting a uniform likelihood function for $\alpha_C > 0.5$ and $\alpha_D < 0.5$. In this case, Firm2's optimal decision is to cooperate under no evidence and E_1 . Whereas, for E_2 , when the probability of AA intervention is close to one, $\mathbb{E}[u(D)|E_2] > \mathbb{E}[u(C)|E_2]$, so Firm2's optimal decision is to defect. These results coincide with those obtained using the TFT strategy shown in the last row of Table 7. Recall that the TFT strategy corresponds to setting $\alpha_C = 1$ and $\alpha_D = 0$ in all **Duopoly** instances.

Now, suppose Firm2 believes that its rival cooperates — with probability $\alpha_C = 0.8$ — if Firm2 cooperates; and cooperates — with probability

Table 7: Firm2's expected utility for different values of α_C and α_D , without evidence, with evidence E_1 and E_2 , for likelihood evidence and for the TFT strategy.

α_C	α_D	without evidence		with evidence E_1		with evidence E_2	
		$\mathbb{E}[u(D)]$	$\mathbb{E}[u(C)]$	$\mathbb{E}[u(D) E_1]$	$\mathbb{E}[u(C) E_1]$	$\mathbb{E}[u(D) E_2]$	$\mathbb{E}[u(C) E_2]$
1	0.25	337	388	329	339	322	298
0.8	0.25	286	316	278	277	271	245
0.6	0.25	238	250	228	219	220	193
0.4	0.25	203	193	190	170	180	152
1	0.2	332	388	326	339	321	298
0.8	0.2	281	316	275	277	270	245
0.6	0.2	231	247	225	217	219	193
0.4	0.2	192	188	183	167	177	149
1	0.1	321	388	321	339	320	298
0.8	0.1	270	316	270	277	269	245
0.6	0.1	219	243	219	215	218	193
0.4	0.1	172	179	171	159	170	143
likelihood		280	313	273	275	268	243
TFT		385	443	390	394	395	353

$\alpha_D = 0.25$ — even if Firm2 defects. This is implemented in the network inserting and propagating evidence `alpha_C=0.8` and `alpha_D=0.25` in each **Duopoly** instance. As we can see in Table 7, Firm2’s expected utility to cooperate, $\mathbb{E}[u(C)] = 316$, is greater than to defect $\mathbb{E}[u(D)] = 286$. Introducing evidence E_1 in **AA_1**, the two decisions become almost utility equivalent. Whereas, under the TFT strategy, E_1 yields an optimal decision to cooperate $\mathbb{E}[u(C)|E_1] = 394$, whereas $\mathbb{E}[u(D)|E_1] = 390$.

Recall that when information about the relevant market is not taken into account, the probability of AA intervention is 0.0189. If the probability that Firm1 cooperates when Firm2 defects is very small ($\alpha_D = 0.1$), then its optimal decision is to cooperate, even for small values of α_C . On the other hand, when $\alpha_D \geq 0.2$, defecting is Firm2’s best choice for $\alpha_C = 0.4$, yielding a different behaviour from that obtained using the TFT strategy. However, using evidence E_1 , when the probability of AA intervention is 0.514, $\mathbb{E}[u(D)|E_1] > \mathbb{E}[u(C)|E_1]$ even when Firm1 is slightly altruistic, $\alpha_D \leq 0.2$ and $\alpha_C \leq 0.6$. Furthermore, if $\alpha_D = 0.25$, then $\mathbb{E}[u(D)|E_1] > \mathbb{E}[u(C)|E_1]$ also for $\alpha_C \leq 0.8$. If the TFT strategy is adopted, Firm2 optimally cooperates both under no evidence and E_1 . Whereas, for E_2 , the associated probability of AA intervention is very large, so that Firm2’s optimal decision is to defect for all values of α_C and α_D considered here.

While the examples shown here are merely illustrative, the number of questions and different strategies that can be analysed is clearly huge and increases with the number of stages considered.

4 Conclusion

We have shown that BN models may enrich the existing methods and results of representing and solving symmetric repeated games. From a game theory perspective, OOBN graphical representation has considerable benefits in comparison to standard game tree representation.

When the antitrust authority starts an investigation, the two potentially merging firms are likely to represent a relevant share of the market, hence they might affect the price of the goods traded. In contrast, the decisions of other firms inside the market, but outside the merged entity, can be assumed

to be irrelevant. In circumstances such as these, a PD duopoly model is a reasonable representation.

From an economic perspective, the methodology we present can be seen as a useful decision support system. It models and integrates the different uncertainty sources deriving from a rival competitor and from the economic environment. Furthermore, the model can be updated as we consider new cases or changes in the market conditions occur.

As is standard in industrial organization the firm is seen as a single decision making unit; generalizations of our OOBN to model firms internal organization could also be considered. Indeed, firm's top and middle management may have different objectives from its owner. An appropriate BN could be built to model these interrelationships and incorporate them into a more general OOBN model. This would yield a more complete and realistic picture of firms' cooperative behaviour. We hope to develop this and other aspects in the future.

Acknowledgement

This research was supported by a PRIN07 Italian Research Grant. We are indebted to Carlo Cazzola, Mauro La Noce and Valerio Ruocco of the Italian Antitrust Authority for providing the data. We thank Mario Tirelli and Filippo Calciano for their helpful suggestions. We also thank Gaetano Bloise for refereeing the paper.

References

- Andreoni, J. and Miller, J. (1993). Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. *Economic Journal*, **103**, 570–85.
- Bangsø, O. and Wullemmin, P. (2000). Object oriented Bayesian networks: A framework for topdown specification of large Bayesian networks and repetitive structures. Technical report, Department of Computer Science, Aalborg University.

- Buntine, W. (1991). Theory refinement on Bayesian networks. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pp. 52–60. P. S. B. D’Ambrosio and P. Bonissone (Eds.), Morgan Kaufmann.
- Cowell, R. G., Dawid, A. P., Lauritzen, S., and Spiegelhalter, D. J. (1999). *Probabilistic Networks and Expert Systems*. Springer Verlag, New York.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of Royal Statistical Society B*, **39**, 1–38.
- Flood, M. (1958). Some experimental games. *Management Science*, **5**, 5–26.
- Friedman, J. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*, **38**, 1–12.
- Jensen, F. V. (2001). *Bayesian Networks and Decision Graphs*. IEEE Computer Society Press, New York.
- Kadane, J. B. and Larkey, P. D. (1982). Subjective probability and the theory of games. *Management Science*, **28**, 113–20.
- Kearns, M., Littman, M., and Singh, S. (2001). Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pp. 253–60. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Koller, D. and Milch, B. (2001). Multi-agent influence diagrams for representing and solving games. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1027–34.
- Koller, D. and Milch, B. (2003). Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, **45**, 181–221.
- Kreps, D. M., Milgrom, P. R., and Wilson, R. (1982). Rational cooperation in the finitely repeated prisoner’s dilemma. *Journal of Economic Theory*, **27**, 245–52.
- Kreps, D. M. and Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, **27**, 253–79.
- La Mura, P. (2000). Game networks. In *In Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 335–42. UAI.
- La Noce, M., Bolasco, S., Allegra, E., Ruocco, V., and Capo, F. M. (2006).

- Merger control in Italy 1995-2003: a statistical study of the enforcement practice by mining the text of Authority resolutions. *International Journal of the Economics of Business*, **13**, 307–34.
- Lauritzen, S. L. and Nilsson, D. (2001). Representing and solving decision problems with limited information. *Management Science*, **47**, 1238–51.
- Roth, A. E. and Murnighan, J. K. (1978). Equilibrium behavior and repeated play of the prisoner’s dilemma. *Journal of Mathematical Psychology*, **17**, 189–98.
- Smith, J. Q. (1996). Plausible Bayesian games. In *Bayesian Statistics 5*, pp. 387–402. J. M. Bernardo and J. O. Berger and A. P. Dawid and A. F. M. Smith (eds), Oxford.
- Steck, H. (2001). Constraint-based structural learning in Bayesian networks using finite data sets. In *PhD Thesis*. Department of Computer Science, University of Munich.