EPSRC
Engineering and Physical Sciences
Research Council

InFoMM
Industrially Focused
Mathematical Modelling

EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling

Bargaining Under Uncertainty

Huining Yang

bp

UNIVERSITY OF
OXFORD

Contents

1 Introduction

Negotiation is an essential interaction in human activities. For example, an oil company such as BP needs to negotiate with foreign governments to gain the rights for the petroleum exploration. It is widely studied in many disciplines, including economics, artificial intelligence, game theory, and social psychology. Of particular interest is a negotiation in which two parties (the buyer and the seller) enter into negotiation rounds and take actions in alternating order to agree on the value of a project. At each round one party proposes an offer and the other party either accepts the offer, or rejects it and proposes a counter-offer. Negotiations stop when both parties agree on the price of the project. Negotiators face two problems: should they accept the offer from the opponent or not? If one party rejects the offer, what is the optimal counter-offer? Negotiators may be trying to optimize various utilities; our interest is finding optimal offers in order to maximize the profit.

The players adopt the time-dependent tactics and can be classified into 3 groups.

A key facet of the negotiation is the *Reservation price* (RP), which is a limit on the price of a good or a service. On the demand side, the RP is the highest price that a buyer is willing to pay; on the supply side, it is the lowest price at which a seller is willing to sell a good or service. If the seller's RP is lower than the buyer's RP, there exists a region of agreement. There are three common tactics [1] : time-dependent tactics, resource-dependent tactics, and behaviour-dependent tactics (Tit for Tat). A typically used time-dependent tactic is built on a decision function $T(t)$ involving two parameters: T_i is the private deadline for player i and $\beta(> 0)$ is the *concession rate*. The concession rate measures how fast one player makes a concession, i.e., accepts an offer from the opponent. Depending on the values of β, there are three types of players: *conceder*, *linear*, and *Boulware*, see Figure 1. When $\beta > 1$, the decision function is concave, indicating that the player will only concede near the deadline, while the conceder has the opposite behaviour.

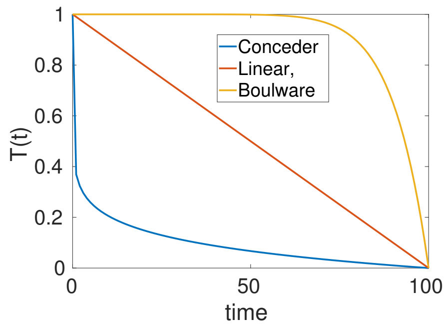

Figure 1 – Three types of players: conceder, linear and Boulware

Each party only knows his own RP and deadline.

We focus on bargaining situations with incomplete information in which each party only knows their own reservation price, deadline, and strategies. This information asymmetry introduces uncertainty in the negotiation process. Other uncertainties may result from the valuation of projects, movement of the market and other unpredicted events.

Our aim is to develop bargaining models that a buyer can use to agree a contract around a single issue – usually price.

Glossary of terms

- **Reservation price (RP) :** A limit on the price of a good or a service.
- **Concession rate:** A parameter measuring the degree of concession.
- **Bayesian learning:** A statistical method where Bayes' theorem is used to estimate the probability of a hypothesis.
- **Non-linear regression:** Data analysis used to model the relationship between two variables.

- **BR model:** A negotiation model based on Bayesian learning and non-linear regression.
- **Discrete Wavelet Transform (DWT) :** A tool used to extract the main trend of data.
- **Gaussian process:** A stochastic process in which every finite collection of random variables has a multivariate normal distribution.
- **GP model:** A negotiation model based on Gaussian processes and DWT.
- **Mapping function:** A function maps a price to some value between 0 and 1.
- **Utility:** A measure of profit.
- **Discounting effect:** The tendency to show a preference for a reward that arrives sooner rather than later.

The BR model is based on Bayesian rule and non-linear regression which is used to predict the opponent's RP and deadline

2 The BR model

In this section, we propose an adaptive model for the buyer, based on Bayesian learning and non-linear regression [2], which can be used to predict the RP and deadline of the opponent at the same time. The key idea is that the buyer uses a simple model for the seller's offer $p_s(t)$, that is built around the seller's initial offer and their unknown RP, which we call RP_s.

Denoting the estimation of the seller's deadline by T_s, we predict the seller's *reservation point* (T_s, RP_s) in two steps: first, we apply non-linear regression to find the correlation between the fitted offers and the seller's offers; secondly, we use the estimate of the correlation with Bayesian rules to update the buyer's belief of the location of the seller's RP and deadline. As the negotiation continues, the buyer receives more offers, thus the prediction will become more accurate. Finally, we calculate the buyer's optimal concession strategy based on the predicted information using a model involving buyer's RP and the buyer's deadline.

Prediction of RP and deadline

In Figure 2, we show a snapshot of a prediction process using the BR model. The x-axis represents the time and the y-axis represents the seller's offer. The red stars are the proposals from the seller. We define a *detecting region* to be a rectangle showing the buyer's estimation of seller's RP_s and T_s. At time $t = 0$, the detecting region represents the buyer's initial guesses of the seller's information. We divide the detecting region into $M = M_1 \times M_2$ equally sized and spaced cells. In order to fit the regression model, we randomly select a point (represented by the blue dot) in each cell which is an estimation of the seller's reservation point. We fit the regression model to the received offers for the seller (red stars) and the random point in cell C_i to obtain an estimation of the seller's concession rate β_s, which gives us the fitted regression line represented by the blue dotted line. We calculate the non-linear correlation between the fitted offers on the regression line and seller's offers, which shows the similarity between the approximation and real data. Next, we apply Bayes' Theorem to update the buyer's belief of the location of the seller's reservation point where the correlation is used as the conditional probability in Bayes' Theorem. We colour the three most likely locations (blocks) in the figure (and, to make the figure clean, we hide the regression lines and random points on these blocks).

Adaptive concession strategy

Having determined an estimate for the seller's reservation point, the next step is to adjust the buyer's bidding strategy adaptively at each round. Depending on the position of the randomly selected point in each cell, we have four scenarios to consider and each one corresponds to a specific optimal concession strategy. We calculate the optimal concession rate β_b^i for each cell and take the weighted average of them to get the buyer's overall concession rate β_b. We use this to calculate the buyer's optimal offers at any given time.

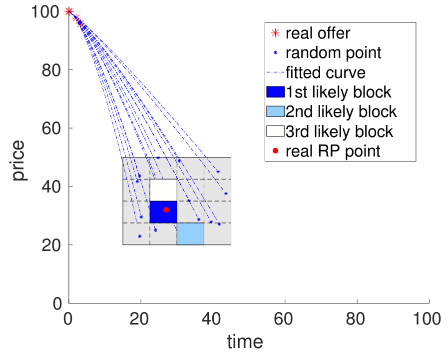

Figure 2 – Snapshot of a prediction process using the BR model.

3 The GP model

The GP model employs DWT and Gaussian processes to model the opponent's behaviour directly.

An alternative way of addressing the negotiation problem is to directly model the trend of the seller's offers and then to search for the optimal price among all possible offers of the opponent in the future. In this section, we discuss another model based on this objective, called the GP model, which involves the DWT technique and a Gaussian process [3, 4].

This model has three key steps: (i) the DWT technique is employed to extract the main trend of the data (offers from the seller), (ii) we fit a Gaussian process to the data to get the predicted mean and variance for the seller's offers in the future, and (iii) we discuss how to generate the optimal counter-offers to reach our target utility.

The DWT technique

The discrete wavelet transform (DWT) is a core technique in signal analysis. When a signal is passed through one layer of a DWT, it is decomposed into 2 parts: a main approximation that represents the main trend of the data, and a detail part which shows the perturbation or noise in the data. We first substitute the seller's offers into a *mapping function* which maps the seller's offers to some values between 0 and 1, which we refer to as the *utility* denoted by U. We then employ a *Daubechies wavelet* [4] with 10 layers to extract the main part of the data.

Gaussian processes

The Gaussian processes are used to find the maximal expected utility of the buyer in the future.

The advantage of employing Gaussian processes to learn the opponent's strategy is that this technique not only gives predictions for the mean of the utility but also provides confidence intervals of the future data, which can be used to show how accurate the predictions are. By fitting a Gaussian process to the utility U calculated through the DWT technique, we obtain the mean and standard deviation of the predicted seller's offers at some future time t between the current time t_c and the deadline T. We note that, in this model, we do not predict the opponent's deadline – this deadline T could be the public deadline of both parties or we could estimate the seller's deadline using the BR model or in some other way. We illustrate the ability to predict the utility using the GP model in Figure 3. The x-axis is the normalized time (between 0 and 1) and the normalized deadline is $T = 1$. The y- axis denotes the utility between 0 and 1 obtained from the mapping function. The red stars indicate the seller's offers scaled between 0 and 1 and the black line shows the main trend in the data found using the DWT technique. We observe that the uncertainty, which is represented by the shaded 95% confidence interval, increases as time approaches the deadline.

Adaptive decision-making process

We assume that the buyer's utility U at a future time t has a normal distribution. Since U always lies in the range $[0, 1]$, we adopt a *truncated* normal distribution, which is used to

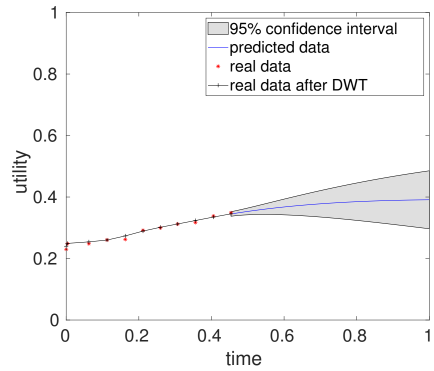

Figure 3 – Snapshot of a prediction process using the GP model.

calculate the expected utility of the buyer. We first consider the optimal time t^*, which is the time (in the future) when the expected utility is maximized. We include a discounting effect and the buyer's risk attitude by involving two parameters in the formula. The maximal expected utility corresponds to the minimum price p_{min} which the seller may propose in the future. We obtain the optimal counter-offer of the buyer by applying time-dependent tactic with the aim of agreeing on p_{min} at the deadline.

4 Results

To measure the effectiveness of our models, we adopt the *linear utility* u_b given by the ratio of the difference between the agreement price and the buyer's RP to the difference between the buyer's initial price and the buyer's RP. This function maps the agreement price into the interval $[0, 1]$ and the higher the linear utility is, the more profit the buyer obtains with their offer. If the final agreement is near the buyer's RP, the linear utility is almost zero. We note that this is a different function to the utility which is obtained from the mapping function in section 3.

In each simulation of the negotiation process, we suppose the seller adopts a time-dependent tactic with a fixed concession rate randomly selected from the interval $[0.5, 2]$. To make the simulation more realistic, we add a negative noise to the seller's offers generated, where the RP for the buyer and seller are randomly drawn from the intervals $[30, 50]$ and $[55, 75]$, respectively. Their deadlines are also randomly chosen from the range $[20, 40]$. In the BR model, we choose to use 36 blocks. In practice, it is often the case that the buyer identifies himself as a conceder, therefore the buyer's concession rate is set to be 0.5 and is fixed throughout the negotiation.

The BR model achieves higher utility than the GP model with seller's offers generated by time-dependent tactic.

By employing the BR model and the GP model, players achieve more efficient outcomes than naive negotiators.

To illustrate the benefit of employing opponent learning models and to compare the outcomes from the BR model against those from the GP model, we assume that the buyer employs one of the two models to learn the seller's behaviour and the seller negotiates 'naively' without any models. We also compare with a completely naive negotiation, in which we assume that both the seller and the buyer select time-dependent tactics and their concession rates are randomly drawn from the range $[0.5, 2]$ and are fixed during the negotiation. We test the effectiveness of two models and the naive negotiations in the presence and absence of noise. We run 500 simulations for each case and we obtain the average linear utility as shown in Figure 4. We observe that using a model gives the buyer a more advantageous price, with the BR model providing a larger utility than the GP model. We show two specific negotiation examples in Figures 5 and 6. In each example, we also plot the offers that a naive buyer would propose, indicated by the black plus. For the BR model, based on the prediction that the seller's deadline is earlier than the buyer's deadline, the buyer's strategy is to wait until the seller proposes their RP near their deadline. For the GP model, the buyer proposes a relatively high price close to the estimated p_{min} and then holds this price until an agreement is reached. In both cases, the buyer obtains more beneficial negotiation outcomes when using one of our models to learn the seller's behaviour than if he goes to a negotiation naively.

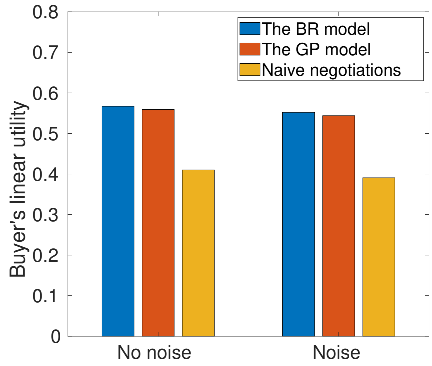

Figure 4 – The average linear utility of two models and naive negotiations against different size of noise.

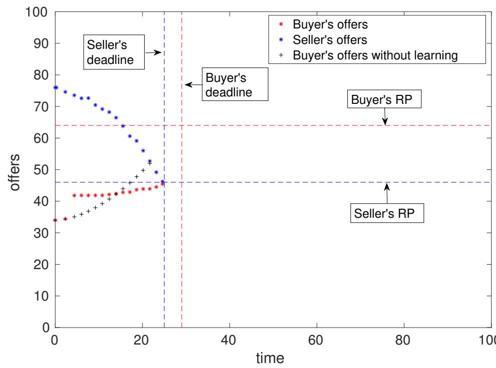

Figure 5 – Negotiation examples applying the BR model.

Furthermore, we test the effectiveness of our two models against seller's offers generated using a behaviour-dependent tactic, which is less often used than the time-dependent tactic in real negotiations. With 100 simulations, the average linear utility of the BR and GP models increases in each case, with the GP model achieving the largest profit. Both methods achieve higher utility in this case since the seller concedes much more quickly than using the time-dependent tactic.

5 Discussion, conclusions & recommendations

We have built and tested models to describe how two parties negotiate a contract, with the aim of helping the buyer to propose optimal offers in a negotiation in order to maximize his profit. In both models, the offers are generated dynamically by analyzing the offers from the seller. We found that the buyer achieves more beneficial negotiation outcomes if they use a model than negotiating naively. The BR model predicts the RP and deadline at the same time and achieves higher utility than that of the GP model with the seller's offer generated using time-dependent tactics. However, the effectiveness may be lower than the GP model if the seller selects another strategy. The GP model works well regardless of the seller's strategy, since we do not impose any assumptions on the seller's tactics. As part of the model, we remove small quantities of noise using the DWT technique, while the deadline is assumed to be public or it needs to be predicted.

We have assumed that only the buyer learns the seller's behaviour, the seller does not adjust his strategy according to the buyer's behaviour. As illustrated in [2], if both parties

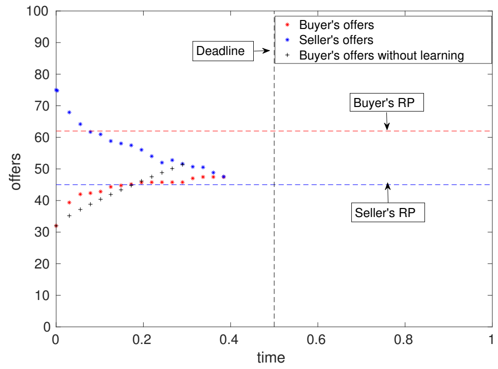

Figure 6 – Negotiation examples applying the GR model.

learn the opponent's strategy using the BR model, the buyer (or the seller) obtains lower utility as both parties are trying to maximize their profits. It would be interesting in future work to explore how the negotiations proceed if the buyer and seller use different models. In real world applications, most negotiations involve multiple issues. To extend the BR model to multi-issue settings, we refer to [5] where the weights of issues are estimated in the first step of the model. Furthermore, both of our models are real-time models which work regardless of whether we have a record of previous negotiations with the counter-party. Other models, in which we need to use such a record to determine the parameters before negotiations start, are likely to provide a much more accurate prediction than the online models.

Potential impact

Paul Barnes, Senior Negotiator for BP said: "*Understanding how a counter-party behaves and reacts in a bargaining situation, particularly where there is asymmetric information available to the parties, is an essential skill for a negotiator to master. In some circumstances however, it's not possible to gain those insights through dialogue. This InfoMM mini-project has opened up a potential mathematical route to help this understanding, and coupled with data science and analytics, there could be real potential in modernizing how negotiators develop their tactics & strategies. Whilst the outcome of the project is not yet at a stage where it could be implemented in practice (our expectation is this requires a multi-year research project), I was very impressed with the level of insight and potential developed in such a short time period, and more so by the dedication and commitment of Huining.*"

References

[1] P. Faratin & C. Sierra & N. R. Jennings, *Negotiation decision functions for autonomous agents*, Robotics and Autonomous Systems, 24(3-4), 159-182, 1998.

[2] C. Yu & F. Ren & M. Zhang, *An adaptive bilateral negotiation model based on Bayesian learning*, Studies in Computational Intelligence, 435 75-93, 2013.

[3] F. Ren & M. Zhang, *Predicting partners' bahaviors in negotiation by using regression analysis*, Z. Zhang, J.H. (eds.) Siekmann KSEM 2007. LNCS (LNAI), vol. 4798, pp. 165-176, Springer, Heidelberg, 2007.

[4] S. Chen & G. Weiss, *An approach to complex agent-based negotiations via effectively modeling unknown opponents*, Expert Systems with Applications 42 2287-2304, 2014.

[5] J. Zhang & F. Ren & M. Zhang, *Bayesian-based preference prediction in bilateral multi-issue negotiation between intelligent agents*, Knowledge-Based Syst. 84(2015) 108-120.