4

MODELLING, SIMULATION AND MACHINE LEARNING

4.1 MODELLING

A model is an abstract description of a subject. Here the subject could be anything from a situation (such as a conflict), a system (e.g. a banking system) to the dynamics of a market (e.g. an auction market). To build a model of a subject means to identify the key components of the subject and describe the relations between them. The hope is to use the model to capture the main behaviour of the subject.

In a model, the components often influence or interact with each other. Such influence or interacting relations could be expressed in any form. They could be expressed mathematically or procedurally, for example, "if component A gets a signal from component B, then A will send signals to components C and D".

Modelling enables one to reason about the subject. We encounter models all the time. For example, at war, army officers put model armies on a map to show their control and influence. This enables them to evaluate moves and counter moves. A war game is a model of real wars. The game SIMS is based on a model of how people interact with each other.

The Basic Alternating-Offers Bargaining model introduced in Section 3.4 is a simple model of bargaining. There the components

DOI: 10.1201/9781003348474-4

are the two players who take turns to make offers and counter-offers to each other. An offer is a percentage of the pie that the player proposes to take. Following are two examples of modelling applied to finance and economics.

Modelling in Interbank Payments

Models have been built to study interbank payment systems. It is a subject studied by central banks around the world, especially after the 2007–2008 financial crisis. When a customer of Bank A pays £1,000 to a customer of Bank B, Bank A must at some point pay £1,000 to Bank B. However, later in the day, another customer of Bank B could be paying a customer of Bank A £800. If the two banks clear their balances by the end of the day, all Bank A needs to do is to pay Bank B the difference, which is £200. The only drawback of doing so for Bank B is that if Bank A goes bankrupt during the day, B will lose £800 which it has already paid its customer. To avoid this risk, the two banks may clear the interbank payments instantaneously. The drawback is that they must maintain a reserve to do so, which eats into their profitability. In the above example, Bank A only needs to use a reserve of £200 to clear the payments by the end of the day, but it must use a reserve of £1,000 were it to clear the balance instantaneously. Central banks have gathered together to design, with the help of models, clearance rules to balance between risk-bearing and reserve burdens.

Modelling in Electricity Markets

Models have been built for designing the rules that govern an electricity market. The electricity market is complex. Multiple suppliers generate electricity to supply end-users through the grid, which distributes electricity to consumers. Excess electricity supply cannot be stored in large quantities and therefore goes to waste. However, a blackout is possible if electricity is under-supplied; the 2000–2001 California blackout was a well-known example. Government regulations and the rules of the electricity market must be designed to ensure a surplus in supply but minimize wastage (costs will eventually be passed on to consumers, so high wastage means high electricity prices). Designing the rules for the electricity market to strike a balance is non-trivial. This has been the subject of research in modelling.

4.2 MODELLING: IMPERFECT BUT USEFUL

Faced with a complicated situation, one often asks: "where should I pay attention to?" Model building helps people to identify the most relevant components and their relations in complicated situations.

In building a model, one is forced to ask what the key components are in the subject, and how these components relate to or interact with each other. For example, in modelling a war situation, the firepower and range of a troop may be quantified. If the modeller believes that the terrain is important, then objects such as rivers, grassland, trees and buildings should be part of the model. If the weather situation is considered to be important, then objects such as rain, snow, wind direction and wind speed should be part of the model too.

Model builders often start with the most basic components and relationships. They knowingly leave out less important components and their relations for a later stage. The initial models are naturally imperfect. A simple model is opted for because it is easier to study. After studying the simple model, more components can be added. More relationships between the components can be added too. An incremental approach enables the modeller to assess the impact of each additional component and relationship.

As more components and relations are added, the model is closer and closer to reality. However, most situations worth studying are complex, hence a model is never a perfect description of reality. We all know that the Basic Alternating-Offer Bargaining model does not describe human bargaining realistically. Communication in human bargaining is a lot more complicated. For example, in a market, a buyer may walk away, hoping that the seller will call him/her back with a better offer. Eye contact and body language are also important in human bargaining which is not in the basic bargaining model.

While a model is never a perfect description of reality, it can still be useful. Building a simple bargaining model helps one to focus on the factors that are most important and study the orders in a subject. When research in the simple bargaining model matures, bargaining theorists may incrementally relax the assumptions or refine the model to make it more realistic.

Due to complexity, one may never be able to build a very realistic model. But what is the alternative? An incremental approach is arguably the only way to study a complex situation. A model is always a simplification of the real situation. All models miss out on something. However, when used properly, models can be useful, as has been demonstrated in many applications.

> ***"All models are wrong, but some are useful".*** **(George Box)**[1]

The topic of modelling will be revisited in Chapter 5 when we discuss portfolio optimization, a financial application.

4.3 SIMULATION: BEYOND MATHEMATICAL ANALYSIS

Models support analysis. When a model is simple, one may be able to mathematically analyse its properties. For the basic bargaining model described in Section 3.4, game theorists have been able to mathematically work out the subgame equilibrium under perfect rationality and perfect information assumptions.

Unfortunately, many interesting models in finance and economics are complicated. For example, in the bargaining model, what if one player does not know the other player's utility decreasing rate? Where would the subgame equilibrium be? That is difficult for mathematical analysis.

When the model is too complicated for mathematical analysis, simulation is often the only reasonable solution. Simulation is used by AlphaGo (explained in Section 2.1) – it runs through millions of moves by the two players in order to evaluate the quality of each possible next move. The unsupervised learning approach to the bargaining problem (explained in Section 3.5) is also a simulation – it runs through possible reasoning by the two players. This kind of simulation is sometimes referred to as Monte Carlo simulation, to reflect the randomness in the process.

Following is a simplified version of how AlphaGo uses simulation to make a move: given a board situation, AlphaGo will generate a random move (to be elaborated below) for the immediate move. Then it generates a subsequent move by the opponent, followed by a subsequent move by the current player, and so on. This simulation brings the game to the end (when all board positions are uncontested), which will tell AlphaGo which side wins. This simulation is repeated millions of times. The immediate move that leads to more wins will be adopted to be the next move.

To improve the efficiency of simulations, AlphaGo does not pick every empty position on the board with equal probability when it generates the next move. More promising positions are given more chances to be picked in the simulation. This is akin to human players spending more time examining the most promising move sequences. Machine learning is used to learn which positions on the board are more promising.

Monte Carlo simulation is not the only way to conduct a simulation. The co-evolution in finding the subgame equilibrium in bargaining, introduced in Section 3.5, can also be seen as a simulation. There one population is used to represent a set of strategies that Player A could adopt, and another population for player B. The two players continually evolve their portfolio of strategies in response to the other player's evolution. Readers are reminded that the main motivation for using co-evolution for the bargaining problem is to relax the perfect rationality assumption in finding subgame equilibrium.

4.4 CASE STUDY: RISK ANALYSIS

After the global financial crisis from 2007 to 2008, the international banking community gathered in Basel, Switzerland, to construct frameworks for securing financial stability in the future.[2] The Basel standards are implemented by individual countries in the form of laws and regulations for their banks. Banks were required to reserve a certain proportion of their capital to protect them to a certain limit in financial crises. The proportion depends on the assets that they hold – a lower reserve is required for assets of lower risks and a higher reserve for assets of higher risks.

The purpose of keeping the reserve is to reduce the chance of bankruptcy by the banks, which would disrupt society. Return on investment is not the regulator's concern. However, from a bank's point of view, the goal is to maximize its return. While it is the bank's duty to comply with the regulatory requirements, maintaining excessive reserves will eat into the bank's profit. In Chapter 5, we shall look closer at the problem of having dual conflicting objectives. In this section, we shall focus on how a bank may satisfy the reserve requirements. The description below is based on the implementation by a technologically advanced financial institute.

As reserves tie up capital, keeping excessive reserves reduces the banks' earning potential. For that reason, banks tend to carry the minimum amount of reserve to meet regulatory requirements. Banks are invited to present evidence to demonstrate that the reserves that they keep meet the requirements. This is where research is required.

Given a portfolio of assets held, a company will have to calculate the minimum reserve that it must hold. To do so, the company models the probability of each asset changing values. For example, if it holds a certain amount of bond X, a model may describe the probability of X losing 0.5% on the next day, the probability of X losing 0.4% the next day, …, etc. How is this model built? It may be based on the historical price changes of X in the past, say, 3 years. Alternatively, the model may be built by using the statistical summary of X's price changes in the past 3 years – the mean and standard deviation, for example.[3] Alternatively, a mathematical model can be built based on the analyst's insight into that asset. For example, the analyst may expect the value of this asset to rise by 0.02% per day, with a standard deviation of 0.01%.

Apart from modelling the price changes in each asset that it holds, the company models the potential change in other external factors, such as the interest rate, inflation rate, etc. Again, changes in these factors can be modelled with historical data, mathematical data or expert insight. The company may also model the dependency between these factors and assets. For example, if the interest rate rises, share prices tend to fall. Such a relationship is nonlinear and complex, but that does not prevent machines from learning from historical data.

Having built these models, the company may start to simulate possible futures. Gains and losses over the next 30 days, say, can be simulated using those models. With these simulations, risk measures can be collected for the portfolio held by the company. For example, with over 100 million simulations, how bad could the portfolio perform in the worst 5% of the simulations? Suppose all these 5% of simulations show a loss of 8% or more, then −8% is called the 5% Value-at-Risk (VaR) of this portfolio, based on the models and simulations. The average loss of the worst 5% is called the 5% Expected Shortfall. These, plus other statistical measures from the simulations, can be presented to the regulators to show the company's reserve meets the regulative requirements.

4.5 ADDING MACHINE LEARNING TO MODELLING AND SIMULATION

We have explained in the preceding section that with models and simulations, a company can assess the risk of a portfolio. We have also explained the use of machine learning in AlphaGo to learn the promise of each position, which makes simulation more efficient. We have explained (in Section 3.5) how machine learning can be used to find the subgame equilibrium in a bargaining problem. In this section, we shall explain the power of adding machine learning to modelling and simulation. As an example, we shall describe how it can be used in designing trading strategies. Following is a cycle for designing trading strategies for a market.

A model–simulate–learn cycle for designing trading strategies:

1. Model the market clearing rules and other traders' behaviour.
2. Model a class of trading strategies, which is the subject of fine-tuning.
3. Simulate the interaction between the subject trading strategies and other traders in the market.
4. Assess the performance of the subject trading strategies.
5. Modify the trading strategies, guided by their observed performances.
6. Run the simulations again.
7. Repeat Steps 4–6 until the results match the desirable results.

To start, the investigator builds a model of the market clearing mechanism. This includes how orders are processed and how they are matched to complete transactions. Stock markets typically adopt a double queue system: the buyers form a queue and the sellers form another queue; traders transact with each other. Foreign exchange markets are typically market-making markets – the market maker sets the buying and selling prices, which they adjust based on supply and demand; traders transact with the market maker.

The modeller may also model trading strategies used by other traders. For example, some traders may use technical rules (they are called technical traders). Some may buy or sell randomly (they are called noise traders) or hold to take profit. As mentioned earlier, models are never perfect, but they could be useful for investigation.

A note on terminology: research that models players and their interactions in a system is sometimes referred to as "agent-based" research, which is a branch of AI. The word "agent" is used to refer to both human traders and algorithmic trading systems.

Next, the investigators must decide how to model a class of trading strategies which could be fine-tuned by machine learning (Step 2). For example, a trading strategy may take the form of a neural network with many internal layers (which is referred to as "deep learning" in AlphaGo, see Section 3.2). Machine learning will then be tasked to adjust the weights on the connections later in the process. The investigator may also decide to maintain a population of neural networks and let them compete against each other.

A trading strategy may also be represented by a tree. A tree is a generic data type which can be used to represent anything computable. Old-fashioned computer scientists were taught that all computer programs can be parsed into trees.[4] In genetic algorithms and genetic programming, which is a population-based machine learning method, a trading strategy is constructed by building blocks. For example, in genetic programming, a trading strategy could be represented by a tree, which is made up of arithmetic operations on financial indicators, such as the previous closing, opening, high and low prices. Trading strategies will be allowed to compete against each other. The poor performers will be eliminated and the fit individuals will be allowed to pass their building blocks on to future generations.

One important point in machine learning is worth reiterating: The choice of representation determines what one can learn. So is the choice of learning method. However, the most important decision is the choice of input and output. A technical trader may input to a learning system technical indicators, such as moving averages. An economist may input into the system fundamental indicators, such as price–earning ratios and macroeconomic indicators, such as interest rates. A poor choice in representation or machine learning method may still produce a mediocre trading strategy for the investigator. But, as explained in Section 3.3, a poor choice in the input variables will limit the ability of the system to produce anything useful.

Once the modelling is complete, the investigators may let the implemented trading strategies interact with each other (Step 3). The performance of each strategy is assessed (Step 4). The investigators must decide what they want to achieve in these traders. Performance could be measured in many ways. For example, it could be measured by the profits made, the percentage of profitable trades or the amount of maximum losses.

How does the system learn? How could it modify the trading strategies based on the feedback (Step 5)? With neural networks, the weights on the network connections record the accumulated knowledge acquired through learning. As in AlphaGo Zero (Section 2.2), this is a form of unsupervised learning. If a population of neural networks is maintained, then networks that performed well so far could be duplicated. Each copy will make minor random modifications to the weights of the connections. The new copies will replace the poor performance in the population. If a decision tree is used to represent trading strategies, a population of decision trees is maintained. Successful strategies are allowed to pass their building blocks to future generations, as explained earlier (see Section 3.5).

After the trading strategies are modified, the simulation will be repeated (Step 6). This simulation and remodelling process can be repeated until the investigator is satisfied with the strategies generated or no improvement is observed in the process.

Above, we have explained how the modelling–simulation–learning cycle could be used to automate the design of trading strategies. There are many other possible approaches. The key point is that the model modification step is laborious if the investigator were to be involved. Machine learning allows the system to test thousands of trading strategies, which would not be practical to do manually.

Some may ask: is efficiency that important? If a trading strategy works, then it is worth spending time to find it. This is a reasonable proposition. It is worthwhile to spend a whole year to find a winning strategy if this strategy reliably makes money in the market for the years to come. Unfortunately, a winning strategy will not be winning if others have found it too. Simple strategies such as the momentum rules and Head and Shoulder pattern probably worked at some point in the past. But as more traders know about them, these known patterns are reflected in the price. To succeed, an investigator must keep looking for new trading strategies. As explained earlier (Section 1.2), the competition is on finding regularities ahead of one's competitors. Therefore, efficiency in inventing new trading strategies matters. By automating the investigation process, the modelling–simulation–learning cycle helps to keep a trader ahead of the others in the game.

4.6 MECHANISM DESIGN

In the previous section, we explained how trading strategies can be designed through a modelling–simulation–machine learning cycle. In this section, we shall look at how rules can be designed for a new market. In economics, this is referred to as "mechanism design", a subject for the Nobel Prize in Economics in 2007.

New markets are created from time to time. We mentioned the electricity market in Section 4.1. This is a continuous market with changing supply and demand. The complexity makes its design challenging. In this market, the electricity demand over time is predicted based on demand patterns in the past. However, the weather, special activities (for example, when a football game is scheduled, electricity demand is expected to surge at half-time because many viewers will put the kettles on simultaneously) and other factors all affect the current demand. Electricity producers must bid the price they are willing to charge and the quantity of electricity that they are willing to produce. To avoid blackouts, the auctioneer (electricity supplier) must buy enough electricity to meet the demands continuously, with demands varying over time. On the other hand, the auctioneer wants the producers to bid the lowest prices under their individual business models. Different suppliers have different capacities. Small suppliers may carry higher production costs, but the auctioneer may not want them to be competed out of the market completely because their presence adds to the stability of the market. Suppliers that generate reusable energy may carry higher production costs too, but the supplier may want to support them for social reasons. In designing these rules, the auctioneer wants to give the producers enough incentive to produce enough electricity. They also want to encourage producers to bid their true valuation of the commodity.

The following is a modelling–simulation–learning cycle for designing rules in a market.

A model–simulate–learn cycle for mechanism design:

1. Model the market rules (which are the subjects of machine learning) and the participants' behaviour.
2. Simulate the interaction between the participants under the market rules.
3. Observe and analyze the results of the simulations.
4. Compare results with desirable results.
5. Modify the market rules in the market model, guided by the observed results.
6. Run the simulations again.
7. Repeat Steps 3–6 until the results match the desirable results.

This process is similar to the design of trading strategies, except in the focus on the models – here the focus is on the market rules (Steps 1 and 5).

Here, the model of the market is a generic framework which may have a set of rules to be selected or not selected, plus a set of parameters to be tuned (Step 1). Machine learning will be used to select those rules and tune the parameters (Step 5). The bidding and acceptance of bids are simulated (Step 2) with results observed and analysed (Step 3).

The auctioneer must decide what they want to achieve in the market. For example, how much blackout risk is it willing to take? That determines how big a surplus (buffer) it must maintain. How much does it want to support more expensive suppliers? That determines how it may accept bids. These objectives must be written down clearly for the modelling–simulation–learning cycle to be automated.

The objective function is used to evaluate the quality of a market model (Step 4). This is crucial to machine learning, which has to decide which models to modify and how to modify them (Step 5).

After the market models are modified, the simulation will be repeated (Step 6). This simulation and re-modelling process can be repeated until the investigator is satisfied with the market rules generated or no improvement (as judged by the objective function) is observed in the process.

What we have described earlier is one way to use a modelling–simulation–learning cycle to automate mechanism design in a market. Without automation, the model modification step could be laborious. Machine learning allows one to test thousands of market models, which would not be practical to do manually.

4.7 CONCLUSION: MODEL-SIMULATE-LEARN, A POWERFUL COMBINATION

Modelling helps one to focus on what to pay attention to. A model helps one to reason about a subject – such as a financial market or a payment system.

Mathematical reasoning is elegant and powerful. But it is only useful for simple situations – such as the Basic Alternating-Offers Bargaining model mentioned in Section 3.4. To study a complex system, simulation is cost-effective; sometimes, it is the only way. We have shown how modelling and simulation enable one to assess the risk of a portfolio.

Adding machine learning to modelling and simulation allows one to find subgame equilibrium in complex game models beyond the Basic Alternating-Offers Bargaining model.

To summarize: modelling, simulation and machine learning could combine to form a powerful tool. Modelling enables simulation. Machine learning helps to improve models. As models are never perfect, such reasoning is always flawed. However, having the means to reason about a subject is better than doing nothing at all.

> ***"More calculation is better than less calculation; some calculation is better than none". (Sun Tze)*[5]**

NOTES

1. G. Box and N. Draper, *Empirical Model-Building and Response Surfaces*, John Wiley & Sons, 1987.
2. Three Basel standards progressively covered more and more aspects of financial aspects. The latest standards were defined by Basel III in November 2010.
3. Provided that the price change distribution follows a Gaussian distribution or other statistical properties.
4. Parsing is no longer a popular part of the computer science syllabus today.
5. Sun Tzu, *The Art of War*, around 5BC ("多算勝，少算不勝，而況於無算乎？").