Using Game AI to Control a Simulated Economy

Paul McCarlie and Aaron Hunter
British Columbia Institute of Technology, 3700 Willingdon Avenue, Burnaby, Canada

Keywords: AI for Games, Rule-based Systems, Machine Learning.

Abstract: We explore the use of Artificial Intelligence (AI) to manipulate a simulated economy. Towards this end, we present work in progress on a macroeconomic simulation that can be controlled by a game player. We view this simulation as a sort of serious game; it can be played as a competition, but it can also be an educational tool through which players learn both about economic principles and the behaviour of AI controllers. The main contribution of this paper is the comparative study of the effectiveness of different AI agents for the manipulation of a simulated environment. Focusing on AI approaches that are common in the gaming industry, we implement four players that use intelligent methods to control the simulation by trying to maximize the economic output. The aim of our work is to illustrate that simple methods from the game AI community can be used to control a complex economic simulation effectively. This work, therefore, supports the common position in the gaming community that simple character-based AI methods can produce competitive game play even for complex tasks. Moreover, we demonstrate that, in the case of this particular simulation, a rule-based reasoner outperforms more sophisticated AI agents.

1 INTRODUCTION

Simulating an economy is an important and challenging problem in many video games. While the economic simulations for games are simpler than those used for real-life forecasting, they are often complex and based on the same economic models. In order to produce Artificial Intelligence (AI) opponents in games of economic simulation, we need to develop tools and frameworks for decision making and manipulation of economic artifacts. In this paper, we present a comparative study of several different AI techniques for controlling a simulated economy. This is work in progress, aimed at determining which AI methods can perform the task effectively. The goal is to create agents that manipulate the economy at a level that is competitive with a human player, using only the basic AI tools that are typically employed in game development. The restriction to common methods from game AI is significant, as we want to explore how effective our AI controllers can be under standard constraints on the development of AI for games. We are also interested in determining which specific AI models produce the strongest controllers.

This paper makes several contributions to existing literature. First, in order to even explore the main problem, we need to develop a realistic economic simulation based on real economic models. We then develop prototype software that implements four different AI controllers for manipulating the economy through simple directions. We illustrate that each approach functions better than a baseline controller, and we determine which approach is the most effective. We argue that these results are immediately applicable in game AI. Our prototype software illustrates which AI technique is most effective at improving the economy, restricting attention to the simple techniques often employed in games. We suggest that, when fully developed, this software will provide a useful testbed for learning about economic principles, as well as learning about the utility of AI for decision making about an economy.

2 BACKGROUND

2.1 Economic Modelling

In this section, we provide a basic introduction to economic modelling. Of course, a detailed discussion of macroeconomic theory is beyond the scope of this paper. We refer the reader to (Ragan, 2020) for a complete introduction.

McCarlie, P. and Hunter, A.
Using Game AI to Control a Simulated Economy.
DOI: 10.5220/0010212306290634
In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence (ICAART 2021) - Volume 2*, pages 629-634
ISBN: 978-989-758-484-8

Our simulation uses the Aggregate Supply-Aggregate Demand (AS-AD) model to determine price level based on supply and demand. This model is used to track taxation, spending, investing and debit. Basically, the price for goods is calculated as an aggregate of *consumption, investment* and *government spending*.

To track long term change in an economy's output, we use the Solow-Swan Growth Model (Donghan et al., 2014). The economic output in this model is determined as an aggregate of *technology, labour, productivity* and *capital stock.* Capital stock change is in turn calculated by a set of different equations, relating population and capital growth.

2.2 Game AI

In this section, we briefly summarize the AI approaches used in our simulation. While game AI borrows techniques from the academic AI literature, there is a distinct approach that requires distinct methods. In particular, the goal in game AI is generally to simulate intelligent characters, rather than to simulate human reasoning; as such there are particular tools and approaches that tend to be implemented and used widely. The AI techniques described in this section are all standard approaches in game development, described in more detail in (Millington, 2019).

The first three approaches are essentially based on rules. First, we use a classic *rule-based system*, in which economic rules based on expert knowledge are defined. These are simple if-then rules, where we keep track of knowledge base that changes as antecedents are triggered. This framework then dictates the actions that the AI should take to grow the economy. The second approach is similar, but it uses *fuzzy logic* in reasoning about the rules. In other words, the rules are based on fuzzy concepts with degrees of truth rather than classical logic (Köse, 2012). The third approach to AI employs *goal-based behaviour,* where the AI has explicit goals and actions rather than simple rules. This approach permits reactive planning in the classical sense of (Georgeff and Lansky, 1987). The advantages of this approach have been discussed in (She and Grogono, 2009).

The last approach to AI used in our simulation is Machine Learning (ML). Specifically, we use regression to learn the relationships between dependent variables based on past information. In this simulation, we actually use a combination of *linear regression, Gaussian regression,* and *Sequential Minimal Optimization regression (SMOreg).* In the present work, we use the Weka framework to implement the ML agent (Frank et al., 2016).

3 ECONOMIC SIMULATION

3.1 Overview

Our software is a game that includes an interactive macroeconomic simulation. The simulation can be controlled by a human player, or by an AI agent. The objective is not only to produce a playable game, but also to provide a useful testing environment for experimenting with economic policies. At the same time, we wanted to determine if standard game AI approaches can operate the simulation more effectively than a human player.

A number of variables are simulated. Broadly speaking, the primary ones are:

- Debt
- Economic Output
- Taxation
- Government spending
- Public investment
- Public consumption
- Technology

The idea of the game is to maximize the economic growth of your economy while keeping debt under control.

There are two separate "debts" simulate in our platform: *government debt* and *public debt*. Government debt is determined by taxation and spending; if players spend more than they are taking in from taxes, government debt must increase. Public debt is a bit more complicated. Public spending depends on the interest rate, which depends on the money supply; a higher money supply means a lower interest rate and therefore more spending. The money supply itself is determined by owned bonds divided by the reserve requirement. In order to increase public spending, the player must either increase the money supply or decrease the reserve requirement. Public income cannot be directly changed; it is determined as a percentage of output. As is to be expected, if spending exceeds income, public debt is incurred.

Debt is automatically paid back periodically, being subtracted from spending, and it must be paid back with interest. The interest is determined by how large the debt is; higher debt means a higher interest rate.

The other important variable is technology. This is what will increase our long term economic output.

It is advanced by public and government spending. This is where the critical tradeoff is: we want to spend enough to maximize investment and long term growth, but not so much that we take on too much debt. If debt gets large enough that it's interest payments cannot be completely paid from spending, and it will directly technological advancement negatively.

3.2 Gameplay

The actual gameplay consists of adjusting policies over game cycles. A game cycle consists of the following:

1. Run the Solow-Swan Growth model.
2. Run the AS-AD model and show the user its indicators.
3. Make policy changes.
4. Run the AS-AD model again and show the user its indicators.

At the beginning of a cycle, the player is presented with the choice of running the simulation manually or with the AI. If they choose manual, then they are shown current economic indicators, and must decide what they wish to change (if anything).

The player has control over 4 different policy levers:

- Taxation
- Spending
- Bonds owned
- Reserve requirements

When prompted how they would like to change each policy, players simply type their responses into the console. When they have made their decisions, the new policies are applied and the player is shown new economic indicators which reflect their changes.

In order to use the AI component, we choose how many cycles we want to simulate. The game then runs as many game cycles as is specified with the AI making the policy change decisions instead of the player.

The program is written in Java and should work on most major Operating systems, dependencies are managed with Gradle.

3.3 Illustrative Example

The game is currently played at the command line. The intention is that this simulation can be a component of a larger game that involves an economy at an underlying component. This is the case, for example, in the Democracy series of games.

When the game is launched, the player is prompted as follows:

Press m for manual play, press a for ai play

If manual play is selected, then a manual cycle starts with a description of both the Solow Model and the AS-AD model:

*-*Solow Model Information*-*
Population Growth rate: 0.0
Total Output: 564.6216173286173

*-*ASAD Model Information pre-adjustment*-*
*-*Output Gap Data*-*
Long Run Aggregate Supply: 564.62161732861

Additional information is displayed, related to aggregate demand, taxation, government spending, inflation and debt. The player is then given options for policy adjustment:

Select option for policy adjustment:
t for taxes
g for government spending
m for money supply
r for reserve requirement

The player can select any option. If they select government spending, for example, they will be prompted as follows:

How much do you wish to change spending by?

Size of spending change needed to close the gap: 14.003562437128632

A number can be entered to increase or decrease government spending, and the simulation will then print out all of the AS-AD model information again. The second line indicates how much change is required to make a difference in the economy; this information is included to help new users. The new model produced after any action will show changes that have occurred as a result of the action taken.

If the player starts off by selecting the AI option, they will be prompted as follows:

Enter number of cycles for AI to run

The given number of cycles will be simulated, using one of the AI models. Results are displayed in the same format as they are for the manual run.

3.4 AI Control

We now briefly describe the implementation of each AI controller, starting with the rule-based reasoner.

3.4.1 Rule-based Reasoner

The rule set used for the rule-based reasoner is simple. We describe the basic rule set here informally. One module of the rule set addresses the situation where the public is rich.

```
(PubBalance > GovBalance)
  ⇒ PublicIsRich
(PublicIsRich & OutputGap > 0)
 ⇒ IncreaseSpending
(PublicIsRich & OutputGap > 0)
 ⇒ DecreaseSpending
(PublicIsRich & OutputGap < 0)
 ⇒ BuyBonds
(PublicIsRich & OutputGap < 0)
 ⇒ DecreaseReserve
```

These rules are applied to determine possible changes. The overall goal here is to reduce the output gap. As such, when there are two different changes are triggered, selection between the two is based on which action will yield the biggest change.

A similar set of rules is included for handling the situation where the public balance is less than the government balance.

3.4.2 Fuzzy Logic

The fuzzy logic module operates similarly to the rule-based reasoner. However, rather than simply comparing the public balance and the government balance, we use fuzzy valued variables to describe the properties of the model. The *fuzzification* of the variable *debt*, for example, is defined in an external file as follows:

```
TERM debt :=
(balanceHighNegative,1.0)
(balanceNeutralNegative,0.0)
```

In other words, we use defined terms such as *HighNegative* and *NeutralNegative* as fuzzy-valued properties. We have corresponding defuzzification definitions for our output variables:

```
TERM surplus :=
  (spendingHighNegative,1)
  (spendingNeutralNegative, 0);
```

These blocks let us assign numeric values to concepts like “high spending” or “neutral spending.” Overall, the fuzzy logic AI differs from the rule-based implementation in that we do not have to hard code increases or decreases in values; the actual numeric values are determined by the fuzzy membership values.

3.4.3 Goal-oriented Behaviour

When Goal-Oriented Behaviour is used, the AI loops through all nine possible policy changes and tests the result for each. The software implements this capability through two functions:

- *tryOption(i):* Returns the complete ASAD model that will be generated by option *i*.
- *getEconomicHealth():* Uses a formula to estimate the economic health, after selecting a particular option.

Hence, the Goal-Oriented behaviour option essentially uses a Markovian approach to make each choice without looking ahead beyond immediate effects.

3.4.4 Regression

The Machine Learning Regression AI uses a classifier learned from past data, using the Weka implementation for regression. The classifier looks at the current state of the AS-AD model, and chooses the policy change that best fits the current state by classifying it with respect to the training data.

4 TESTING

4.1 Simulation Testing

Before discussing the performance of the game and the AI player, it is important to note that the simulation itself was extensively tested to ensure that it worked correctly. In other words, each action that a player can take was tested to determine if the simulated economy changed in the expected manner. Consider, for example, increasing taxation. The expected outcomes of increasing taxation include the following:

- Equilibrium output should decrease
- Government Balance and Total Government Balance should increase
- Inflation Rate and Average Inflation Rate should decrease
- Price Level should decrease

All of these changes were validated in our testing. Similarly, it was verified the expected changes occur for all of the other changes a player can make in the system.

4.2 Preliminary User Testing for Manual Gameplay

A small collection of four players tested manual gameplay with an early version of the software. Several superficial changes to the software were made to address their concerns.

- The current block-based layout of the AS-AD display was developed in response to readability concerns.
- The output gap was included in the AS-AD model display for different parameters, as it helps users determine how much they should change property values.
- Cycle numbers were added to the output.
- The AS-AD model and the Solow model are actually both displayed in sequence, as users found it confusing when they had to choose between the two models.

Overall, our users found the final interface presented here to be understandable and usable for gameplay. Of course, a larger user study is required for further improvement and validation.

4.3 AI Testing

Each AI algorithm was tested to determine how well it manages the simulated economy. The focus of this testing is on the Long Run Aggregate Supply (LRAS). This is an indicator that essentially measures the long-term economic output; when running the simulation, a high LRAS score indicates strong economic performance.

For comparison, we first ran a baseline test, in which 10 cycles run without any manipulation of the economy. We then ran the same simulation for 10 cycles with a human player, as well as each AI algorithm. The results are given in Table 1.

Table 1: Comparing Player Performance.

Player	Final LRAS
Baseline	784
Human	899
Rule-based Reasoner	892
Fuzzy Logic	879
Goal-Oriented Behaviour	860
Machine Learning	834

Before analysing the results, we emphasize this is a simple test of a prototype system. More detailed testing and refinement would be beneficial. Nevertheless, several interesting observations can be made about the test results:

- All of the AI players outperform the baseline; it appears that manipulations made are beneficial.
- All of the AI players perform worse than a human with game experience.
- The best AI player is the rule-based version, while the worst is the machine learning version.

The results here are somewhat disappointing in terms of the machine learning approach, as we expected it to be more effective. Instead, the most effective AI follows very simple rules that essentially encode human knowledge. However, on the whole, the resuls are encouraging. It is easy to bankrupt the economy through poor play, but this did not happen. All of our AI players were better than a null agent, suggesting that they each have some value as opponents.

5 DISCUSSION

5.1 Economic Simulation

The simulation used in this prototype is unique from the perspective of gaming. We are not aware of any games that allow the user to “play” as the central bank. We are also not aware of any other game where the Solow Model and the AS-AD model are combined into a single macroeconomic simulation. The result is a complex model, with many variables that interact in a manner that is difficult to predict. As such, we suggest that this simulation actually provides an interesting framework that can be used in games involving commerce.

Hence, from the perspective of pure gameplay, we argue that our prototype has been successful. We also remark that, while our focus has been on treating the simulation as a game, it can also be used as a tool for studying real economies. As such, it is possible to actually experiment and learn about effective economic policies through this simple game.

5.2 Game AI

In terms of AI, our focus here has really been on experimenting with different kinds of AI players. Informally, we had two main goals:

- The AI players should be based on simple techniques that are common in the game AI community.
- The AI players must perform better than a static or random player.

We were able to achieve both of these goals and implement four different approaches for experimentation.

In terms of the comparative results, our work is consistent with the standard perspective for AI in games. While sophisticated AI based on machine learning is incredibly useful for many practical problems, it is often the case that these methods are not as beneficial in creating interesting AI opponents. In this paper, we have seen that the best AI is also the simplest: a player that uses a simple rule-based system.

Hence, even when a game is based on a complex simulation, our work suggests that one need not employ complex AI to create competitive and believable opponents.

5.3 Future Work

There are several directions for future work. Clearly, one direction is pure software development. While our simulation is interesting to study in the abstract, it is not currently a very interesting game. In order to improve this situation, it needs to be packaged as a framework that can be implemented in a more stimulating game environment.

Ideally, our simulation will be delivered as a component for games that involve gameplay beyond economic simulation. Hence, the economy underlying a complex simulation game will be simulated in a realistic manner and controlled by an appropriate artificial agent. We are specifically interested in packaging our simulation as a component of educational software, in which people can use the simulation to learn about the nature of economic decision making and manipulation.

In terms of the AI, we intend to implement and test additional approaches. For example, we have not included any AI methods based on rigorous Knowledge Representation formalisms, nor have we included any neural-network based Machine Learning. It would be valuable to test such methods as part of a complete implementation. Moreover, we intend provide more detailed testing with a range of economies. At present, we believe that the poor performance of the ML agent is partially due to the paucity of starting data. We would therefore like to provide more a detailed comparison with more initial data, and more varied test cases. We leave this detailed comparative analysis for future work.

6 CONCLUSION

This paper is a report on work in progress, focused on using AI agents to control a simulated ecnomy in a game setting. We have described a realistic economic simulation that can be controlled as a playable game. The player of the game acts as the central bank (or government) and takes actions to try and improve the economic output. We have defined four different AI algorithms for playing the game, based on standard AI methods used in the gaming community. While the AI players do not outperform a skilled human at present, the AI players do make decisions that improve economic output when compared to a simple baseline controller.

Overall, this work demonstrates that simple game AI methods can effectively control a simulated economy. The comparison between different AI methods is still at an early stage, but preliminary results suggest that a simple rule-based system provides better performance than more complex methods, including an approach based on machine learning.

REFERENCES

Donghan, C., Hui, Y., and Longfei, G., "A Generalized Solow-Swan Model", *Abstract and Applied Analysis*, 2014.

Frank, E., Hall, M. and Witten, I. The WEKA Workbench. Online Appendix for "Data Mining: Practical Machine Learning Tools and Techniques", Morgan Kaufmann, Fourth Edition, 2016.

Georgeff, M. and Lansky, A. Reactive reasoning and planning. In *Proceedings of AAAI*, 677–682, 1987.

Köse, U. Developing a fuzzy logic based game system. *Computer Technology and Application*. 3, 510-517, 2012.

Millington, I. *AI for Games*. CRC Press, 2019.

Ragan, C. *Macroeconomics*. Pearson Canada, 2020.

She, Y. and Grogono, P. A procedural planning system for goal oriented agents in games. In *Proceedings of the 22nd Canadian Conference on Artificial Intelligence*, 245–248, 2009.

Winterstein, D. “A Simple Intro to Gaussian Processes (a Great Data Modelling Tool).” *Platypus Innovation: A Simple Intro to Gaussian*, Platypus Innovation, 2016.