# Learning to Forecast and Forgetting the Past: A Simulation of Individual Evolutionary Learning with Unknown Asset Returns *

**K. Laughren**†

December 20, 2020

**Abstract**

This paper reports the design and preliminary results of a computer simulation of individual evolutionary learning (IEL) in a fantasy sports decision making context. Agents use strategies that are conditionally active to form beliefs about uncertain asset returns. These strategies are subjected to a genetic algorithm (GA) that involves reproduction, crossover, mutation, and selection. A strategy's fitness is determined by its complexity and prediction error when active. Sensitivity analyses demonstrate that (i) use of a selection operator in the GA leads to long-term improvement in forecast accuracy, (ii) agents evolve strategies with lower prediction error if they face a non-zero cost of complexity, (iii) successful strategies make disproportionate use of future-looking information relative to historical information, and (iv) more accurate beliefs lead to higher profits.

*Working Paper - DO NOT CITE OR DISTRIBUTE

†Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, BC V5A 1S6, Canada, e-mail: klaughre@sfu.ca

# 1 Introduction to fantasy sports

In 2015, there were estimated to be over 55 million people playing fantasy sports in North America, with approximately 70% of them participating in paid leagues[1].

A paid fantasy sports league is a collection of agents competing over a real-money prize to be awarded at the end of a professional sports league season. Agents build a portfolio of assets (a team of athletes from a professional league) and based on the statistical outcomes those players achieve, the agent is awarded points. The agent who receives the most points over the course of a season wins a monetary prize, funded by a buy-in payment from participating agents.

## 1.1 Translating fantasy sports to economics

In every fantasy sports context, it is reasonable to assume that all pertinent information is publicly available, including the prices and expected return of an asset in a given period (see Section 5 for a visual sample of the decision screen from Yahoo! Sports). It is also reasonable to assume that all decisions are incentive compatible, as agents compete for real-money prizes. Individuals gather information and make observable portfolio decisions on a publicly available website, and so we can apply panel methods to their observed decision making.

[1] Fantasy Sports Trade Association, "Industry Demographics"; accessed 19 September 2017; http://fsta.org/research/industry-demographics/

Subsequent to fine-tuning this pilot simulation, my goal is to scale the data and choice problem in the simulation to exactly match that which is faced by real users of fantasy sports and see if this IEL simulation can replicate real-life observed decision making in this context. This pilot study's primary goal however is to demonstrate the application of computational learning (via inductive reasoning) to this problem setting and identify any interesting themes about the characteristics of the most effective strategies and parameterizations that result in the fastest or most effective learning on the part of agents.

### 1.1.1 Daily Fantasy Sports (DFS): Repeated, Static, Non-Strategic Decisions

Each day in which there are professional matches, all individuals face the same portfolio choice problem. Using identical information, budgets, and prices, they build a portfolio of assets from a finite number of choices. Each asset faces an uncertain payoff of points which is directly driven by the statistical outcomes of one professional athlete that day. Financial prizes are paid to the individual who builds the highest paying portfolio, with prizes for the top score each day and occasionally also for the highest aggregate payoff over a professional sports season.

In this context, there is no scarcity, and so it translates to a decision theory context. For example, there is no limit on the number of individuals with the "Tom Brady" asset in their portfolio.

The decision is also static, as each individual starts each day with a new budget and an empty portfolio, so there is no effect of previous decisions on present decisions[2].

The decision context remains the same from observation to the next, however, the prices and expected return of each asset change as a result of changes in the predicted outcomes of the athletes the assets represent. So we could think of this as *repeated observation of decision making under uncertainty*.

**This daily fantasy sports (DFS) context - with its repeated, non-strategic decisions - is the one applicable to the IEL simulation that is the focus of this paper.**

There is a second prominent style of fantasy sports known as a Rotisserie fantasy sports league. Inductive reasoning is less applicable to the Rotisserie context because Rotisserie decisions are strategic in nature (actions of one agent affect the payoff of others). The strategic nature of Rotisserie is a result of assets being rival in that environment; only one agent per contest may have the asset "Tom Brady" in his portfolio. The remainder of this paper is relevant only to the DFS context.

[2]Except in cases of an aggregate payoff at year end

## 1.2 Relation to existing literature

See the appendix for discussion of related literature. The simulation design and results presented here are preliminary and so it is not crucial for the reader at this time to understand the scientific contribution of the paper. Briefly, this paper uses a conditional belief strategy set up as in Arthur *et al.* (1997), and evolves those strategies through an IEL process, with the steps of the GA closely mirroring that of Hommes and Lux (2013). This is explained in much greater detail in the next section.

# 2 A model of fantasy sports belief formation

## 2.1 Primitives:

### 2.1.1 Environment Notation and Parameters

| | | |
|---|---|---|
| $J$ | | integer number of decision making agents |
| $j$ | $\in \{1, ..., J\}$ | indexes a specific agent |
| $M$ | | integer number of strategies held by each agent |
| $m$ | $\in \{1, ..., M\}$ | indexes a specific strategy |
| $T$ | | integer valued number of decision rounds |
| $t$ | $\in \{1, ..., T\}$ | indexes a decision round |

### 2.1.2 Portfolio Notation and Parameters

| | | |
|---|---|---|
| $N$ | $= 20$ | integer number of available assets (finite) |
| $i$ | $\in \{1, ..., N\}$ | indexes an asset |
| $C$ | $= 5$ | integer number of asset categories |
| $c$ | $\in \{1, ..., C\}$ | indexes a category |
| $n_c$ | $= 4 \quad \forall c$ | integer number of available assets in category c |
| $\sum_{c=1}^{C} n_c = N$ | | categories partition asset space |
| $k_c$ | $= 2 \quad \forall c$ | number of assets to be chosen for category C |
| $\prod_{c=1}^{C} \binom{n_c}{k_c}$ | | $=$ 7,776 total number of possible portfolios |
| $\pi_{i,t}$ | | random payoff of asset $i$ in round $t$ |
| $p(i,t)$ | | exogenous, common knowledge price of $i$ at $t$ |
| $W$ | | finite budget constraint on portfolio choice |
| $W \leq \sum_{i=1}^{N} p(i,t)\mathbb{1}\{i \text{ chosen by } j \text{ at } t\}$ | | $j$'s budget constraint |

### 2.1.3 State, Strategy and Belief Notation and Parameters

| | | |
|---|---|---|
| $state\_length = 13$ | | length of the bitstring representing an asset state |
| $condition\_length = 13$ | | length of the conditional part of strategy |
| $belief\_length = 7$ | | length of the bitstring representing how to form beliefs when a strategy is active |
| $X(i,t)$ | $\in [0, \infty)$ | default ex ante belief of $\pi_{i,t}$ (common knowledge) |
| $Active(j,m,i,t)$ | $= 1$ | if $j$'s strategy $m$ is active for asset $i$ at $t$, |
| | $= 0$ | otherwise |
| $\mu^j_m(i,t)$ | $\in [0, \infty)$ | $j$'s belief of $\pi_{i,t}$ under strategy $m$ |
| $\mu^j_m(i,t)$ | $= X(i,t)$ | if$Active(j,m,i,t) = 0$ |
| $\mu^j_m(i,t)$ | | follows random process specified by 7 belief bits if$Active(j,m,i,t) = 1$ |
| $Belief^j(i,t) = \frac{1}{M}\sum_{m=1}^{M} \mu^j_m(i,t)$ | | $j$'s mean belief of$\pi_{i,t}$ |

## 2.2 State variables and strategies

### 2.2.1 State variables

Each asset $i$ has an accompanying state variable for each $t$. This state variable is a 13-digit binary string. The bits represent both historical and future-looking information about the asset as follows:

- Bits 1 and 2: WHAT IS THE PROBABILITY AN ASSET HAS RETURN ZERO?
  This is public injury info offered by the league underlying the dgp (e.g., NFL).
  = (0,0) if there is an ex ante prob = 0 the asset generates a 0 return this period
  = (1,1) if there is an ex ante prob = 1 the asset generates a 0 return this period
  = (0,1) represents a 1/3 probability of 0 return
  = (1,0) represents a 2/3 probability of 0 return

- Bits 3 and 4 - DID IT BEAT EXPECTATIONS? i.e. Payoff / Expectation for t-1
  = (0,0) if at t-1 (Payoff / expectation ) $\leq 0.8$
  = (0,1) if at t-1 $0.8 <$ (Payoff / expectation ) $\leq 1$

= (1,0) if at t-1 1.0 < (Payoff / expectation ) ≤ 1.2
= (1,1) if at t-1 1.2 < (Payoff / expectation )

- Bits 5 and 6 - Past Payoff/ Expectation for t-2 i.e. Bits 3 and 4 lagged one period

- Bits 7 and 8 - WAS IT VALUABLE? i.e. Payoff / Price for t-1
  = (0,0) if at t-1 (Payoff / Price) ≤ 0.8
  = (0,1) if at t-1 0.8 < (Payoff / Price ) ≤ 1
  = (1,0) if at t-1 1.0 < (Payoff / Price ) ≤ 1.2
  = (1,1) if at t-1 1.2 < (Payoff / Price )

- Bits 9 and 10 - Payoff/ Price for t-2 i.e. Bits 7 and 8 lagged one period

- Bits 11, 12, 13 - ARE PAYOFFS INCREASING? i.e. Payoff(t-1) minus Payoff (t-2)
  Bit 11 = 1 if Payoff(t-1) > Payoff (t-2), = 0 otherwise
  Bit 12 = 1 if Payoff(t-2) > Payoff (t-3), = 0 otherwise
  Bit 13 = 1 if Payoff(t-3) > Payoff (t-4), = 0 otherwise

### 2.2.2 Strategies

Each strategy combines a 13 digit "condition" string and a 7 digit "belief" string. The bits in the condition string take values $\in \{0, 1, \#\}$ and determine whether a strategy is "active" by comparing the condition string to the state variable of each asset (also 13 bits). This is analogous to the setup in Arthur *et al.* (1997). However, the particular meaning of each bit is different in my setting, as is the type of beliefs that can be formed. In Arthur *et al.* (1997) the beliefs were restricted to a linear form, whereas in this setting they are formed based on anchoring and adjustment, with an independent probability that the belief is truncated to zero.

**Condition String**

The condition string determines whether a strategy is "active" for a given asset, i.e., it indicates whether a strategy is used to form beliefs for that asset.

Strategy $m$ is active for asset $i$ in period $t$ if every bit in the condition string of $m$ matches the corresponding bit in $i$'s state variable bitstring.

Condition bits can take values from $\{0, 1, \#\}$. If a condition bit is a #, then the strategy is indifferent to whether the asset state has a 1 or 0 in that position. If a condition bit is a 0 or a 1, however, then it must match the asset state bit in that position, otherwise the strategy is deactivated. For example, if strategy $m$'s 3rd condition bit is a 0 and asset $i$'s 3rd state bit is a 1, then strategy $m$ is not active in forming beliefs for asset $i$, regardless of whether any or all other bits match.

**Belief String**

When the condition string of a strategy $m$ is active for an asset $i$, then the 7 digits in the belief string determine the distribution from which an agent's beliefs are drawn. All agent beliefs are anchored on an exogenous, common knowledge expectation of an asset's return for time $t$. That common belief is then augmented independently for each active strategy of agent $j$.

- Bits 14, 15, 16: WHAT IS THE PROBABILITY AN ASSET HAS RETURN ZERO?

  = (0,0,0) if the agent assigns prob = 0 the asset generates a 0 return this period

= (0,0,1) represents a 1/7 probability of 0 return
...
= (1,0,1) represents a 5/7 probability of 0 return
...
= (1,1,1) if the agent assigns prob = 1 the asset generates a 0 return this period

- Bits 17:20 - What is the distribution of the random additive error
  = (0,0,0,0) a uniform RV is drawn with mean $-X_{i,t}$, i.e., $\sim U[-2X_{i,t}, 0]$
  = (0,0,0,1) a uniform RV is drawn with mean $-\frac{13}{15}X_{i,t}$, i.e., $\sim U[-\frac{28}{15}X_{i,t}, \frac{2}{15}X_{i,t}]$
  = (0,0,1,0) a uniform RV is drawn with mean $-\frac{11}{15}X_{i,t}$, i.e., $\sim U[-\frac{26}{15}X_{i,t}, \frac{4}{15}X_{i,t}]$
  ...
  = (1,0,1,0)a uniform RV is drawn with mean $\frac{9}{15}X_{i,t}$, i.e., $\sim U[-\frac{10}{15}X_{i,t}, \frac{20}{15}X_{i,t}]$
  ...
  = (1,1,1,0)a uniform RV is drawn with mean $\frac{13}{15}X_{i,t}$, i.e., $\sim U[-\frac{2}{15}X_{i,t}, \frac{28}{15}X_{i,t}]$
  = (1,1,1,1) a uniform RV is drawn with mean $X_{i,t}$, i.e., $\sim U[0, 2X_{i,t}]$

## 2.3 Individual Evolutionary Learning

### 2.3.1 Strategy and State Initialization Parameters

In period t = 0, agents are randomly seeded with strategies, and state_bits 3:13 - which depend on non-existent historical information - are randomly generated using the following parameters:

| | |
|---|---|
| $p_{\#} = 0.8$ | prob *condition_bits* 1:13 = #, i.e., are indifferent to state |
| $p_1 = 0.1$ | prob *condition_bits* 1:13 =1, i.e., are active only if state bit =1 |
| $p_0 = 1 - p_{\#} - p_1$ | prob *condition_bits* 1:13 = 0, i.e., are active only if state bit =0 |
| $p_{zero} = 0.25$ | prob *belief_bits* 14:17 take value = 1 |
| $p_{error} = 0.5$ | prob *belief_bits* 18:20 take value = 1 |
| $p_{init} = 0.1$ | prob *state_bits* 3:13 take value = 1 at t = 0 |

### 2.3.2 Strategy Fitness Specification and Parameters

A strategy's fitness is a linear function of its mean squared error when it is active (MSE) and its COMPLEXITY, which takes a value from {0,1,...,13} and counts the number of condition bits which are not #s.

$$fit(j,m,t) = -\theta \sum_{i=1}^{N} (\mu_m^j(i,t) - \pi(i,t))^2 - \tau COMPLEXITY(j,m)$$

| | |
|---|---|
| $\theta = 1$ | fitness penalty of MSE is fixed = 1 |
| $\tau \in [0,\infty)$ | the cost of COMPLEXITY relative to MSE varies by simulation |

For ease of notation I will sometimes write:

$MSE(j,m,t) = \frac{1}{N}\sum_{i=1}^{N}(\mu_m^j(i,t) - \pi(i,t))^2$, and

$meanMSE(t) = \frac{1}{J}\frac{1}{M}\sum_{j=1}^{J}\sum_{m=1}^{M} MSE(j,m,t)$

### 2.3.3 Genetic Algorithm Specification and Parameters

At the beginning of each $t \geq 1$, strategies go through a genetic algorithm (GA). In all simulations, the GA contains the steps *reproduction, crossover,* and *mutation.* In a subset of simulations, a fourth step - *selection* - is performed after *crossover* and *mutation* to choose the Parent strategy or the Child strategy that resulted from crossover and mutation based on their $fitness(j,m,t-1)$. So Parent strategies are evaluated on their actual fitness from the previous period, while Child strategies are evaluated on the same data a Parent strategies, but this fitness is "hypothetical" in that Child strategies were never actually used to form beliefs at $t-1$.

The parameters affecting the GA are as follows:

$$
\begin{aligned}
p_{cross} &= 0.1 && \text{prob. a pair of strategies has a crossover occur} \\
p_{mute} &= 0.03 && \text{prob. that a bit in a binary bitstring is mutated} \\
p_{mute3} &= 0.03 && \text{prob. that a bit in a \{0,1,\#\} bitstring is mutated} \\
p_{mute3,\#} &= 0.8 && \text{cond. prob. that a mutating 0 or 1 bit becomes a \#} \\
p_{mute3,1} &= 0.5 && \text{cond. prob. that a mutationg \# bit becomes a 1}
\end{aligned}
$$

### Reproduction

Reproduction takes strategies from $t-1$ and brings them into use for an agent at $t$ with a probability proportional to its fitness. Mechanically, the strategy $\hat{m}$ with the lowest fitness for $j$ is identified and assigned a value of 1. Then all other fitness are scaled to this value and the reproduction sample is drawn according to a roulette wheel style draw:

$$
\begin{aligned}
\hat{m} =& argmin_m \quad fit(j,m,t-1) && \text{find lowest fitness strategy} \\
pos\_fit(j,m,t-1) =& fit(j,m,t-1) - fit(j,\hat{m},t-1) + 1 && \text{normalization} \\
\rightarrow pos\_fit(j,\hat{m},t-1) =& 1 && \text{non-zero prob } \forall m \\
repr\_prob(j,m,t) =& \frac{pos\_fit(j,m,t-1)}{\sum_m pos\_fit(j,m,t-1)} && \text{roulette wheel reproduction}
\end{aligned}
$$

The resulting strategies for $j$ are denoted $SParent(j,m,t)$.

### Crossover

Crossover randomly pairs without replacement each strategy $Sparent(j,m,t)$ that is an output of the reproduction process. Then with probability $p_{cross}$ these two strategies are subject to a crossover. When this occurs, a random integer between 1 and the strategy length is drawn, then all bits after this integer from strategy A are appended to strategy B, and vice versa.

### Mutation

After crossover, mutation randomly occurs on the resulting strategies. Specifically, each of bits 1:13 is mutated with probability $p_{mute3}$ and each of bits 14:20 is mutated with probability $p_{mute}$. A mutated bit always changes its value. If one of bits 1:13 is a 0 or 1 and is selected for mutation, it becomes a # with probability $p_{mute3,\#}$, i.e., a 0 becomes a # with probability $p_{mute3,\#} = 0.8$ or becomes a 1 with probability 0.2.

### Selection

After crossover and mutation, the resulting strategies are denoted as a set $SChild$. The fitness of each strategy $SChild(j, m, t)$ is compared to the corresponding strategy from the pre-existing set $SParent(j, m, t)$, and the one with the higher fitness is carried into the subsequent belief formation and decision making period for the agent. The fitnesses are based on their respective performances at $t-1$, i.e., a belief $\mu^j_{\tilde{m}}(i, t-1)$ is formed for every $\tilde{m} \in SChild$ and used to construct $fit(j, \tilde{m}, t-1)$ which is compared to $fit(j, m, t-1)$ where $m \in SParent$ is the parent of $\tilde{m}$ prior to crossover and mutation.

## 2.4 Portfolio Selection

Agents use $Belief^j(i,t)$ to develop a belief for each of the 7,776 possible portfolios. A grid search is used for an agent to identify the portfolio they expect to yield the highest profits. If that portfolio is not affordable, i.e., the budget constraint $W \leq \sum_{i=1}^{N} p(i,t)\mathbb{1}\{i \text{ chosen by } j \text{ at } t\}$ is violated, then that portfolio is dropped from consideration and the remaining portfolio with the highest expected return given $Belief^j(i,t)$ is chosen. This feasibility loop is conducted until the agent's best affordable portfolio is identified.

## 2.5 Exogenous data

For $i = 1, ..., 20$ and for $t = 0, ..., 16$ I collect the following exogenous data from Yahoo! Sports from the 2016 NFL season in the accompanying file *2016 NFL FSPORTS DATA SAMPLE20X17.xlsm*, which loads into the accompanying python simulation *fsports vSubmit.py* using the package *pandas*:

| | |
|---|---|
| $X(i,t) \in [0,\infty)$ | Exp return of $i$ at $t$, common knowledge, exogenous |
| $p(i,t) \in [0,\infty)$ | Price of $i$ at $t$, common knowledge, exogenous |
| $Inj(i,t) \in \{0,1,2,3\}$ | Common knowledge $prob(\pi(i,t) = 0) = Inj(i,t)/3$ |
| $State(i,t)[0,1] \in \{0,1\}X\{0,1\}$ | Inj(i,t) maps first 2 bits of asset $i$ state |
| $\pi(i,t) \in [0,\infty)$ | realized profit of $i$ at $t$, unknown to agents at $t$ |

### 2.5.1 Repeating data

With an NFL season limited to 17 decisions, there are not many iterations of the GA to converge on a set of strategies. To extend the exogenous dataset of prices, expectations, and profits, the dataset $\{\{X(i,t), p(i,t), Inj(i,t), \pi(i,t)\}_{i=1}^{20}\}_{t=1}^{17}$ is repeated 10 times, i.e., agents make 170 decisions in all the simulations presented, and $X(i,2) = X(i,19) = X(i,36)$.... So an agent at $t = 19$ faces the exact same information as an agent at $t = 2$ but his strategies have had 17 periods of evolutionary learning relative to the strategies he used at $t = 2$.

# 3 Sample Simulation Dynamics

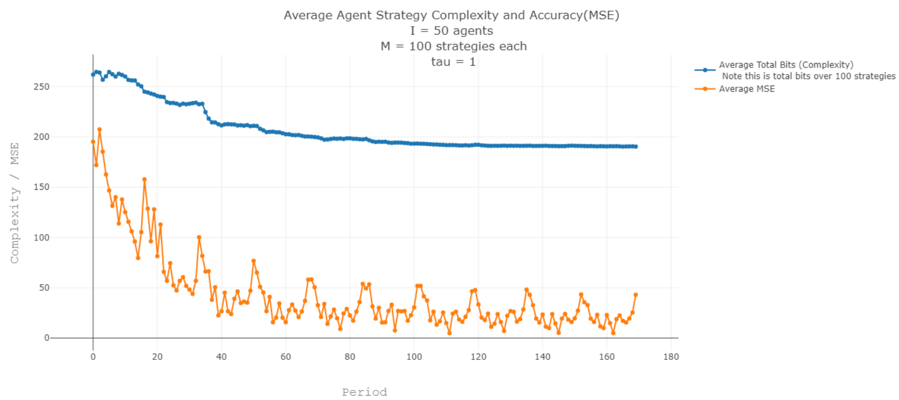


Figure 1: **Mean Complexity and MSE of strategies over 170 periods**. Selection = 1, i.e., selection operator is active. Other parameters included in figure. The mean MSE of all periods is 44.15. Note the decline in MSE over time; linear regression of MSE on round t yields a slope of -0.864, which provides a measure of the rate of learning. Average total complexity at t = 170 is 190 bits or 1.90 non-# condition bits per strategy

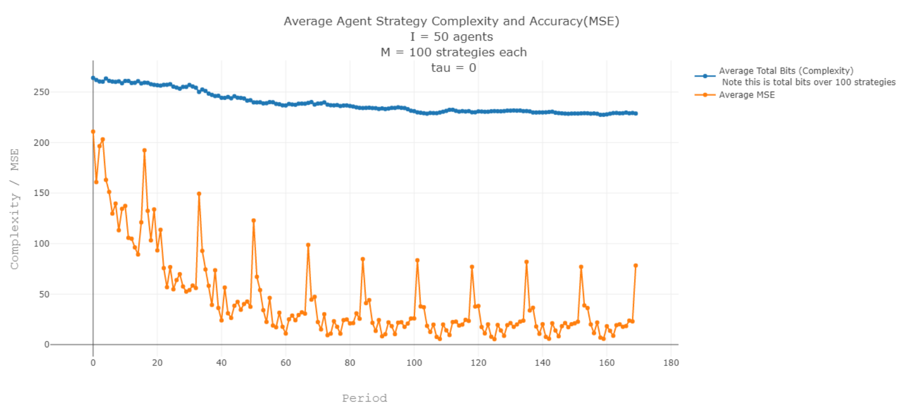


Figure 2: **Mean Complexity and MSE of strategies over 170 periods**. Selection = 1, i.e., selection operator is active. Mean MSE = 45.54. MSE slope (learning) is -0.754. Final average 2.29 non-# condition bits per strategy

Figures 1 and 2 show simulations where the selection operator is active. Figures 3 and 4 below use the same parameters but no selection operator, and it is obvious that MSE does not decline as quickly and strategies become more complex over time, as opposed to less complex over time when Selection is active.

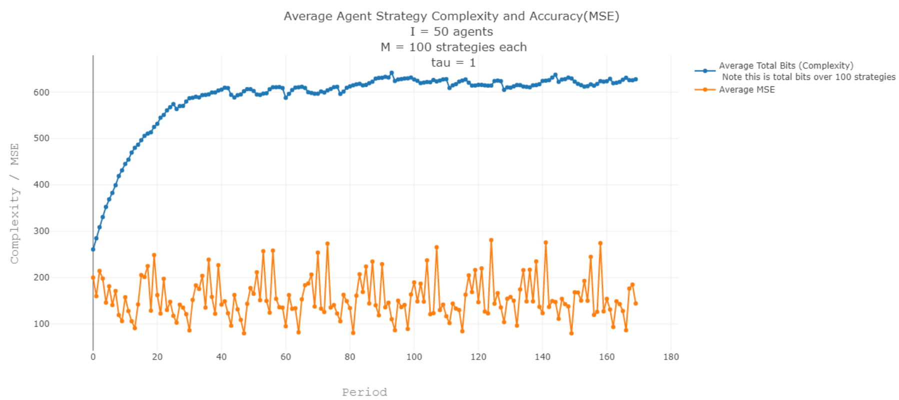


Figure 3: **Mean Complexity and MSE of strategies over 170 periods**. Selection = 0, i.e., selection operator is NOT active. Mean MSE = 155.68. MSE slope (learning) is -0.010. Final average 6.28 non-# condition bits per strategy.**NOTE: scale differs from Figures 1 and 2 because of greater complexity**

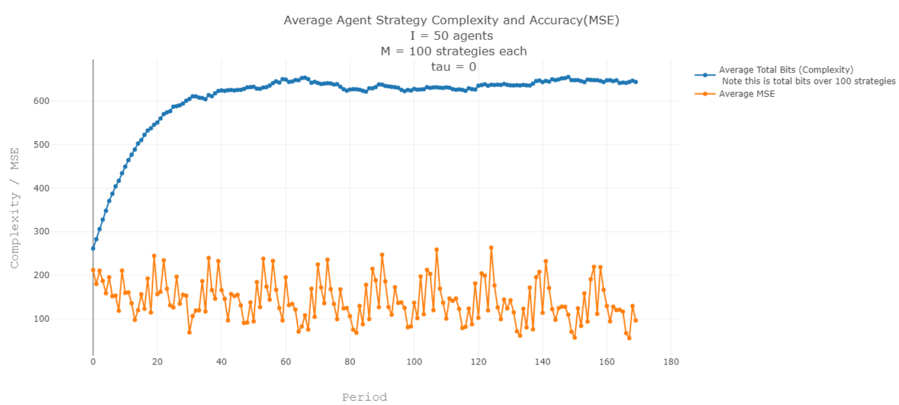


Figure 4: **Mean Complexity and MSE of strategies over 170 periods**. Selection = 0, i.e., selection operator is NOT active. Mean MSE = 142.35. MSE slope (learning) is -0.242. Final average 6.43 non-# condition bits per strategy

# 4 Early Findings and Sensitivity Analyses

## 4.1 The benefit of a complexity cost

**Table A: MSE and learning sensitivity to tau, Selection = 1**

| I agents | M strategies | T periods | tau complex cost | Selection 1 if Selection | meanMSE units | Learning slope of MSE |
|---|---|---|---|---|---|---|
| 50 | 100 | 170 | 0 | 1 | 45.54 | -0.754 |
| 50 | 100 | 170 | 1 | 1 | 44.15 | -0.864 |
| 50 | 100 | 170 | 2 | 1 | 48.97 | -0.786 |
| 50 | 100 | 170 | 5 | 1 | 48.95 | -0.806 |
| 50 | 100 | 170 | 10 | 1 | 51.66 | -0.859 |
| 20 | 200 | 170 | 20 | 1 | 53.79 | -0.820 |
| 50 | 100 | 170 | 50 | 1 | 66.24 | -0.898 |

*Note MSE is lower and the rate of change (learning) is faster when tau = 1 vs tau = 0*

**Table B: MSE and learning sensitivity to tau, Selection = 0**

| I agents | M strategies | T periods | tau complex cost | Selection 1 if Selection | meanMSE units | Learning slope of MSE |
|---|---|---|---|---|---|---|
| 20 | 200 | 170 | 0 | 0 | 156.42 | -0.299 |
| 20 | 200 | 170 | 1 | 0 | 123.04 | -0.285 |
| 20 | 200 | 170 | 20 | 0 | 152.98 | -0.309 |
| 20 | 200 | 170 | 50 | 0 | 155.29 | -0.409 |
| 20 | 200 | 170 | 100 | 0 | 165.08 | -0.297 |
| 20 | 200 | 170 | 200 | 0 | 178.17 | -0.003 |

*Note MSE is lower when tau = 1 versus tau = 0*

## 4.2 The benefit of a selection operator

**Table C: MSE and learning sensitivity to Selection**

| I agents | M strategies | T periods | tau complex cost | Selection 1 if Selection | meanMSE units | Learning slope of MSE |
|---|---|---|---|---|---|---|
| 50 | 100 | 170 | 0 | 1 | 45.54 | -0.754 |
| 50 | 100 | 170 | 0 | 0 | 142.35 | -0.242 |
| 50 | 100 | 170 | 1 | 1 | 44.15 | -0.864 |
| 50 | 100 | 170 | 1 | 0 | 155.69 | -0.010 |

## 4.3 The (lack of) benefit of a larger set of strategies

**Table D: MSE and learning sensitivity to M, various controls**

| I agents | M strategies | T periods | tau complex cost | Selection 1 if Selection | meanMSE units | Learning slope of MSE |
|---|---|---|---|---|---|---|
| 50 | 100 | 170 | 0 | 1 | 45.54 | -0.754 |
| 50 | 200 | 170 | 0 | 1 | 47.14 | -0.772 |
| 50 | 100 | 170 | 1 | 1 | 44.15 | -0.864 |
| 50 | 200 | 170 | 1 | 1 | 46.92 | -0.818 |
| 20 | 100 | 170 | 50 | 0 | 142.78 | -0.314 |
| 20 | 200 | 170 | 50 | 0 | 155.29 | -0.409 |
| 20 | 300 | 170 | 50 | 0 | 163.29 | -0.229 |

## 4.4 The benefit of conditioning on future-oriented information

In an asset's bitstring representing the state, only the first 2/13 bits carry future oriented information. Correspondingly, only the first 2/13 bits in a strategy condition string are able to condition on future-oriented information. Also, because strategy bits 1:13 are all randomly initialized according to the same parameters $p_{\#}$ and $p_1$, the bits in initial strategies are consistently made up of $\approx \frac{2}{13} \approx 15.4\%$ "future bits", that is, bits which condition a strategy on the future information.

Figures 5 and 6 below demonstrate that over time, a greater proportion of strategy condition bits are those that condition on future information. For example in Figure 5 when tau=1, the average complexity at $t = 170$ is 165.7 or 1.66 bits per strategy, of which 0.96 are "future bits". That is, the strategies evolve from being 15% percent future-oriented to being $\frac{96}{166} \approx 58\%$ future oriented. In Figure 6 with tau = 0 we see a similar pattern, although future-oriented bits increase only to 44%, likely because the zero complexity cost allows less useful non-future bits to hang around for more periods.

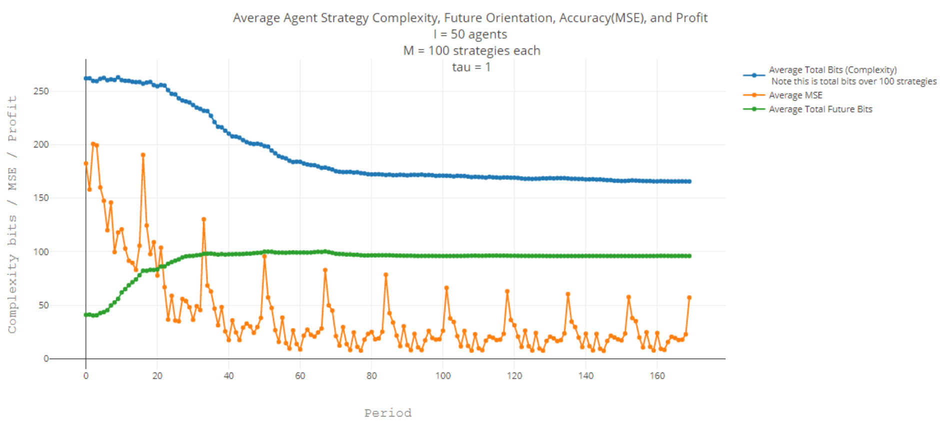


Figure 5: **Mean Complexity, Future Orientation and MSE of strategies over 170 periods**. Selection = 1, i.e., selection operator is active. Mean MSE = 40.59. MSE slope (learning) is -0.773. Final average 1.66 non-# condition bits per strategy, **of which, 58% are future-oriented**

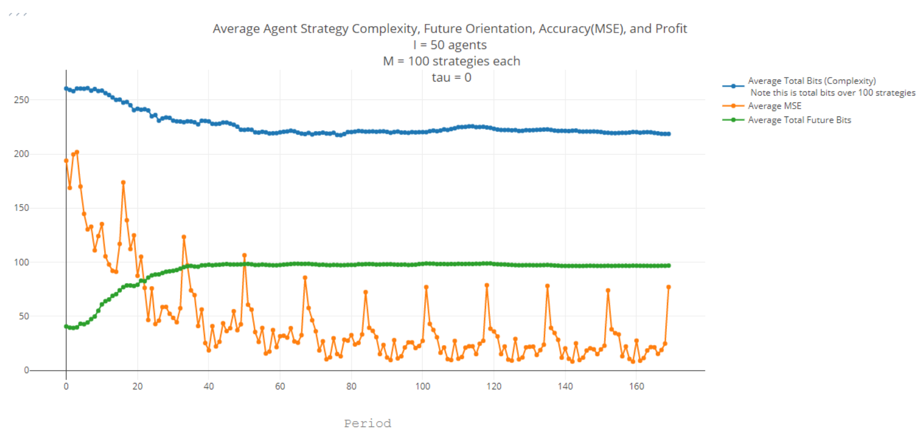


Figure 6: **Mean Complexity, Future Orientation and MSE of strategies over 170 periods**. Selection = 1, i.e., selection operator is active. Mean MSE = 45.35. MSE slope (learning) is -0.790. Final average 2.19 non-# condition bits per strategy, **of which, 44% are future-oriented**

My hypothesis is that it is more valuable to condition on future-oriented information in this setting because much of the pertinent historical information is included in the common knowledge expectation $X(i,t)$, while future oriented information seems to be excluded. For an example of such information being excluded, see the first blue box in Figure 9. In that case, the asset "Adrian Peterson" is on IR, indicating the asset is certain to return a zero profit. However he still has a positive value of $X(i,t)$ so there is clear value for agents who condition on the future-oriented information to update away from the anchor point expectation $X(i,t)$ to a value closer to zero.

# 5 Preliminary Results: Are accurate beliefs profitable?

Among the treatment results demonstrated above, the most stark difference is the long-term accuracy improvement when the GA includes a selection operator. Fixing all of the other parameters at their most accurate levels, Figures 7 and 8 below demonstrate the difference in simulation results between one with selection and one without, including the profits of the portfolios that follow from the beliefs agents form. In neither situation is there substantial improvements to profit over time, however the mean profit with selection is 130.3, whereas with no selection it is lower at 123.3. Further, Figure 8 demonstrates that without a selection operator, the proportion of conditional strategy bits which are future oriented stays fixed near $\frac{2}{13}$ where it is initialized, whereas the simulation with selection migrates to condition heavily on future information (as discussed above). Thus, there is a profit incentive to improve belief accuracy in this way over time, however it is disappointing that there is not more of an upward trend over time.

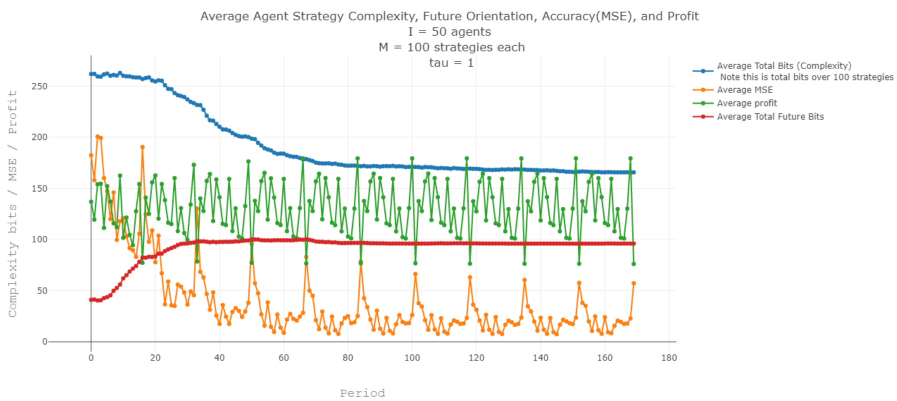


Figure 7: **Mean Complexity, Future Orientation, MSE, and Profit of strategies over 170 periods**. **Selection = 1, i.e., selection operator is active.** Mean MSE = 40.59. MSE slope (learning) is -0.773. Final average 1.66 non-# condition bits per strategy, of which, 58% are future-oriented. **Mean Profit = 130.31**

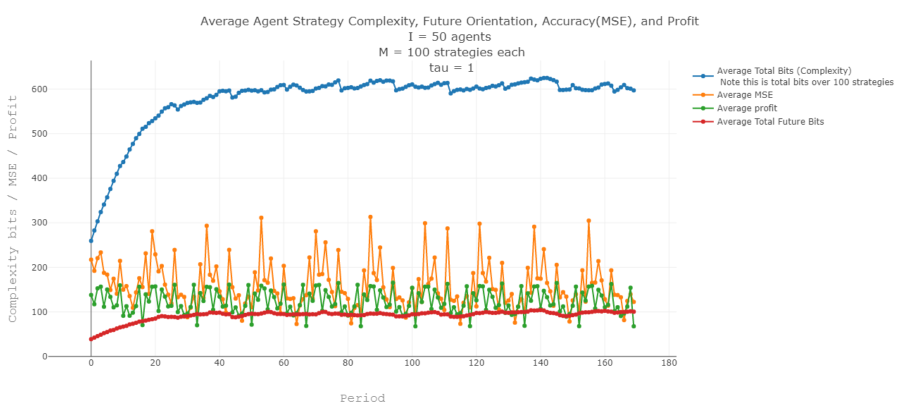


Figure 8: **Mean Complexity, Future Orientation, MSE, and Profit of strategies over 170 periods**. **Selection = 0, i.e., selection operator is NOT active**. Mean MSE = 159.04. MSE slope (learning) is -0.118. Final average 5.97 non-# condition bits per strategy, of which, 17% are future-oriented. **Mean Profit = 123.31**

# 6 Next Steps

I believe the results presented here suggest that this pilot study should be scaled to beyond N=20 assets so that it can be used to predict belief formation in the identical context to that faced by fantasy sports players in the field. Then I can create summary measures to describe field data and see if I can generate similar behaviour through calibration of the simulation, and in doing so discover something about learning, belief formation, and profitable strategies in this context.

There are three significant difficulties to having the necessary pieces to accomplish such a study:

1. A non-grid-search method of choosing a portfolio for an agent given his beliefs. Grid search works for the case presented with 7,776 possible portfolios, and can even extend to the low millions, but generates a memory error for bigger possibility sets like the ones faced by agents in the field. As such, it may make sense to implement another separate GA that allows agents to inductively select a portfolio given their beliefs without checking every possible one and whether it is affordable.

2. Gathering substantial data on assets, e.g., expanding the exogenous data set of 2016 NFL data from $N = 20$ to $N \approx 400$ potential selections

3. Gathering substantial data on decisions made by individuals in real life fantasy sports, e.g., the proportion of the active population of human players which chose asset $i$ and/or which chose portfolio $q$.

Problem 1) above is the most open-ended, but it is not critical to move forward with the project in the near term. The portfolio selection process is the slowest part of implementing the coded simulation, but the process of forming beliefs and updating strategies based on their past accuracy and complexity can be studied without addressing the portfolio problem. Long term, this will need to be addressed to match simulations to field data, as we can only see the portfolios of real subjects, not their underlying beliefs.

Problem 2) can be addressed through brute force copy and paste using a web browser and a spreadsheet, but it would be preferable to develop a python loop which queries the fantasy sports provider, e.g., sending multiple queries to Yahoo! Sports using their proprietary query language YQL.

Problem 3) requires the type of brute-force-or-clever-code approach of problem 2, but presents one additional problem. That is, I need to develop summary measures of field behaviour that is directly comparable to simulation data, similarly to the measures (m1,m2,m3) developed by Arifovic and Ledyard (2017, unpublished) to categorize the choices of experimental and simulation observations into "Nash", "Alternate" or "Other". This classification is likely built on a proportion of subjects choosing a particular asset or portfolio but requires further thought and reading.

I look forward to comments from any interested readers!

# 7 Sample Real-World Analog of Decision Environment

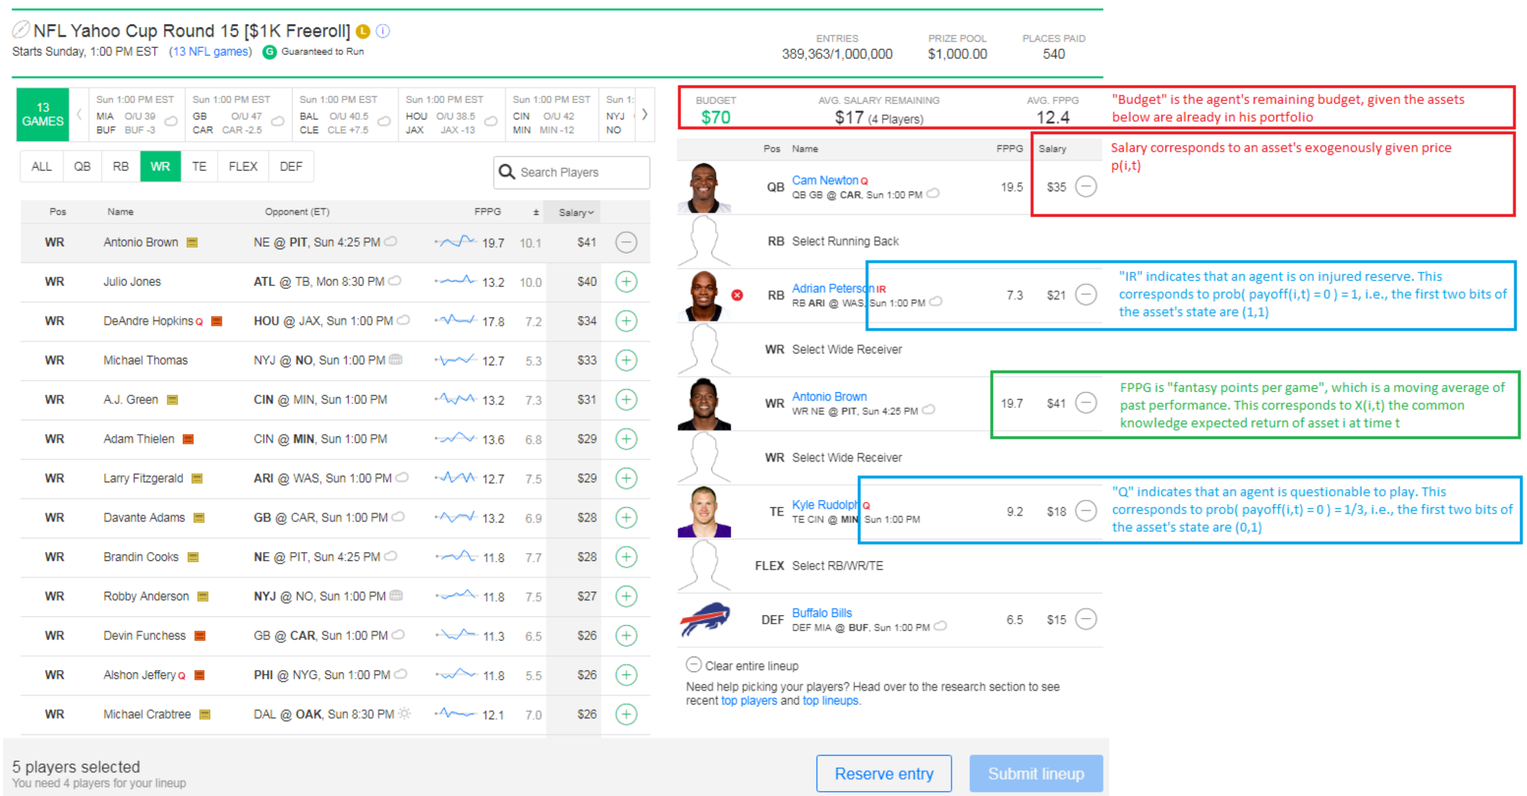


Figure 9: **A Daily Fantasy Choice Interface - NFL week 15 2017**. I have indicated in color above where an agent can see the relevant information and how it ties in to the notation expressed above.

## References

ANUFRIEV, M. and HOMMES, C. (2012). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal.Microeconomics*, **4** (4), 35–64, copyright - Copyright American Economic Association Nov 2012; Document feature - Equations; ; Last updated - 2013-03-25.

ARIFOVIC, J. (1994). Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, **18** (1), 3 – 28, special Issue on Computer Science and Economics.

— and LEDYARD, J. (2007). Call market book information and efficiency. *Journal of Economic Dynamics and Control*, **31** (6), 1971 – 2000, tenth Workshop on Economic Heterogeneous Interacting Agents.

— and MASCHEK, M. K. (2006). Revisiting individual evolutionary learning in the cobweb model – an illustration of the virtual spite-effect. *Computational Economics*, **28** (4), 333–354.

ARTHUR, W. B., HOLLAND, J. H., LEBARON, B., PALMER, R. G. and TAYLER, P. (1997). Asset pricing under endogenous expectations in an artificial stock market. *Available at SSRN, url: https://ssrn.com/abstract=2252*.

AXELROD, R. (1997). *The evolution of strategies in the iterated prisoner's dilemma*. Cambridge University Press.

Bacon, P. M. and Moffatt, P. G. (2012). Mortgage choice as a natural field experiment on choice under risk. *Journal of Money, Credit and Banking*, **44** (7), 1401–1426.

Bondt, W. F. and Thaler, R. (1985). Does the stock market overreact? *The Journal of finance*, **40** (3), 793–805.

Boudreau, J. W. and Shunda, N. (2016). Sequential auctions with budget constraints: Evidence from fantasy basketball auction drafts. *Journal of Behavioral and Experimental Economics*, **62** (Supplement C), 8 – 22.

Cabot, L. V., Anthony N. Csoka (2006-2007). Fantasy sports: One form of mainstream wagering in the united states gaming law symposium. *John Marshall Law Review*, **40**, 1195.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T. H. and Olsen, T. (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *The Quarterly Journal of Economics*, **129** (3), 1141–1219.

Ehrman, N. J. (2015). Out of bounds: A legal analysis of pay-to-play daily fantasy sports. *Sports Lawyers Journal*, **22**, 79.

Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, **112** (2), 631–645.

Haigh, M. S. and List, J. A. (2005). Do professional traders exhibit myopic loss aversion? an experimental analysis. *The Journal of Finance*, **60** (1), 523–534.

Heffetz, O., O'Donoghue, T. and Schneider, H. S. (2016). Forgetting and heterogeneity in task delay: Evidence from new york city parking-ticket recipients. *NBER Working Paper No. 23012*.

Hommes, C. and Lux, T. (2013). Individual expectations and aggregate behavior in learning-to-forecast experiments. *Macroeconomic Dynamics*, **17** (2), 373401.

Karcher, R. T. (2006-2007). The use of players' identities in fantasy sports leagues: Developing workable standards for right of publicity claims. *Penn State Law Review*, **111**, 557.

Karg, A. J. and McDonald, H. (2011). Fantasy sport participation as a complement to traditional sport consumption. *Sport Management Review*, **14** (4), 327 – 346.

Knetsch, J. L. (1989). The endowment effect and evidence of nonreversible indifference curves. *The American Economic Review*, **79** (5), 1277–1284.

Lee, W.-Y., Kwak, D. H., Lim, C., Pedersen, P. M. and Miloch, K. S. (2011). Effects of personality and gender on fantasy sports game participation: The moderating role of perceived knowledge. *Journal of Gambling Studies*, **27** (3), 427–441.

List, J. A. (2001). Do explicit warnings eliminate the hypothetical bias in elicitation procedures? evidence from field auctions for sportscards. *The American Economic Review*, **91** (5), 1498–1507.

— (2003). Does market experience eliminate market anomalies?*. *The Quarterly Journal of Economics*, **118** (1), 41.

— (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, **72** (2), 615–625.

Martin, R. J. and Nelson, S. (2014). Fantasy sports, real money: Exploration of the relationship between fantasy sports participation and gambling-related problems. *Addictive Behaviors*, **39** (10), 1377 – 1382.

Nesbit, T. M. and King, K. A. (2010a). The impact of fantasy football participation on nfl attendance. *Atlantic Economic Journal*, **38** (1), 95–108.

— and — (2010b). The impact of fantasy sports on television viewership. *Journal of Media Economics*, **23** (1), 24–41.

Randle, Q. and Nyland, R. (2008). Participation in internet fantasy sports leagues and mass media use. *Journal of Website Promotion*, **3** (3-4), 143–152.

Suh, Y. I., Lim, C., Kwak, D. H. and Pedersen, P. M. (2010). Examining the psychological factors associated with involvement in fantasy sports: an analysis of participants' motivations and constraints. *International Journal of Sport Management, Recreation and Tourism*, **5**, 1–28.

Vriend, N. J. (2000). An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, **24** (1), 1 – 19.

# 8 Appendix

## 8.1 Relation to literature

### 8.1.1 Literature with related context and questions

My simulation and the study by Hommes and Lux (2013) are both interested in agents learning to forecast, and my simulations GA on strategies uses the same four steps - Reproduction, Crossover, Mutation, and (S)election. However a key difference between the two studies is that profits are exogenously determined by sports outcomes in my study, whereas profits are endogenously determined by the price that results from the set of submitted forecasts by agents in their paper.

This potential for endogenous feedback between expectations and profits also exists in the learning to forecast study by Anufriev and Hommes (2012). My study further differs from Anufriev and Hommes (2012) because my strategies are formed endogenously from a random initialization, whereas those authors initialize agents with one of a small set of heuristics beliefs (e.g., adaptive, weak trend following, strong trend following, etc).

### 8.1.2 Literature with related methodologies

This study is similar to the IEL setting in Vriend (2000), in that agents cannot observe other agent's strategies or beliefs, nor can they sample them as in Vriend's social learning setting from the same paper. I do have multiple agents in my simulations but that is to induce some stability to the reported results[3]. There is no ability to learn from others in my setting and so the parameter $J$ is really just like a sample size.

Like Arifovic and Maschek (2006), my simulation allows for the calculation of hypothetical payoffs in evaluating a strategy. Specifically, the set of strategies SChild that results from crossover and mutation has its hypothetical fitness calculated for the previous periods data, and if the selection operator is on, the strategies in SChild are only adopted if they are fitter than their parents.

This simulation is similar to Arifovic and Ledyard (2007) in that it seeks to use positive methods to answer a normative question, namely "what are effective strategies for belief formation in fantasy sports?". My study is also comparable to this one in that I also seek to compare (field) experiment data to the simulation in the long term.

In both my study and that of Arifovic and Ledyard (2017, unpublished), we seek to simulate agents learning to use a known, efficient strategy. In their case it was the strategy to alternate in a repeated Battle of Sexes game. In my case it is learning to condition on future oriented information which re-

[3]exploiting a WLLN

veals something about the probability of payoff truncated to zero.

Finally, I can say that my simulation is similar to the paper on evolving new strategies in the iterated prisoner's dilemma by Axelrod (1997). In both situations, agents are playing against a fixed environment; my study has agents using exogenous information and receiving exogenous payoffs, while Axelrod has agents evolving strategies against a fixed set of strategies submitted by other researchers. In either case, the learning we observe is only relevant to the fixed setting, however it is more robustly observed than the learning that relies on endogenous feedback on beliefs in other learning to forecast experiments that requires sensitive calibration of parameters.

## 8.2 Relation to other literature

### 8.2.1 Literature specific to a fantasy sports context

Boudreau and Shunda (2016) is a rare example of a methodology similar to my own, treating fantasy sports as a set of incentive compatible decision data from the field and using it to study more broad or abstract economic questions. These authors are particularly interested in bidding behaviour observed in rotisserie and head-to-head competitions in which agents participate in an auction for scarce assets prior to competing. They demonstrate a distinct pattern early overbidding followed by underbidding. To a lesser extent, Karg and McDonald (2011), who study the complementary nature of fantasy sports to traditional sports consumption, also apply economic analyses to fantasy sports data.

There are also many empirical and correlational studies of fantasy sports participation. For example there are studies of the effect of fantasy sports participation on television viewership, game attendance, and mass media use (Nesbit and King, 2010b,a; Randle and Nyland, 2008) as well as studies of the correlation between participation and other forms of gambling (Martin and Nelson, 2014). There are also studies where participation is the dependent variable, which study the effect of such covariates as personality and gender (Lee *et al.*, 2011), and psychological factors such as motivations and constraints (Suh *et al.*, 2010).

Finally there is a largely unrelated legal literature on the status of fantasy sports as a method of gambling. This includes (Cabot, 2006-2007; Karcher, 2006-2007; Ehrman, 2015) . Such literature would only be informative to my use of its data for field study if there is reason to believe that data from one jurisdiction is the result of illegal behaviour that is legal elsewhere, and as a result the agents participating in the illegal jurisdiction are structurally more risk-loving, all else equal.

### 8.2.2 Field experiment literature (relevant for longer term project)

When using field data collected from individual choices made for online leagues, fantasy sports can continue the type of economic research pioneered by John A. List, who advanced the field of experimental economics by attempting to replicate theoretical predictions and experimental regularities from the lab with subjects in the field who make decisions with their own money in a familiar context (List (2001), List (2004)). These studies allowed researchers to answer questions of external validity with greater nuance, for example List (2003) found that market experience with collectibles *reduces* the anomalous endowment effect first demonstrated in the lab by Knetsch (1989), whereas Haigh and List (2005) demonstrated that experienced commodity traders at the Chicago Board of Trade demonstrated a *higher* degree of the economic anomaly of interest - myopic loss aversion - demonstrated in the lab by Gneezy and Potters (1997). Developing a database of fantasy sports decisions will allow researchers to compare their experimental findings between subject pools and between familiar and abstract contexts in such economic fields as learning, conflict, and portfolio choice to better understand the robustness of behavioural anomalies across contexts. For example, does the empirical regularity of overreaction to idiosyncratic bad news in the stock market found by Bondt and Thaler (1985) also apply to the market of fantasy sports assets?

Another, more recent economic literature to which this proposal relates is that on behavioural field experiments, which has exploited large datasets of individual choices in real life contexts. One aim of this literature is to test the out-of-sample predictive power of behavioural models, such as the work by Heffetz *et al.* (2016) which successfully predicts individual levels of procrastination in paying a parking ticket, and uses treatment variation in the information received by drivers in the study to demonstrate that this behaviour is better explained with a model of *forgetting* than one of *rational inattention*. Another aim of this literature is to intricately understand the welfare effects of policies at an individual rather than an aggregate level, for example Chetty *et al.* (2014) studied the effect of changing retirement savings subsidies on Danish individual finances, and found that only those already saving enough for retirement are affected by the marginal policy change, suggesting that analyzing a policy on aggregate variables will misrepresent its true welfare effects. Supposing individuals demonstrate overreaction to short-term results in fantasy sports as they do in the stock market, is it driven by an identifiable subset of individuals, and is there an information context in which this problem is attenuated (e.g., with different timing/content of information)? In this way, the fantasy sports project might help to both identify and evaluate policy tools aimed at reducing the welfare costs of behavioural anomalies, particularly those which deal with decision making under uncertainty and contests.

In a more general sense, this proposal is concerned with observable choices of individuals in the field when payoffs are risky or uncertain. Bacon and Moffatt (2012) observe mortgage choices in the field (e.g., between fixed and variable rate, and different terms), and use the data to study individual risk preferences. This paper also uses methods applied in economic learning models, particularly in implementing a genetic algorithm to study how individuals learn - that is, change their behaviour in light of additional information or experience. In particular, this paper borrows genetic algorithm methodology from both Arifovic (1994) and Arthur *et al.* (1997) to model evolutionary learning over time.